


SECOND EDITION

MATLAB for Behavioral Scientists



David A. Rosenbaum,
Jonathan Vaughan, & Brad Wyble

MATLAB

for Behavioral Scientists

Second Edition

Written specifically for those with no prior programming experience and minimal quantitative training, this accessible text walks behavioral science students and researchers through the process of programming using MATLAB. The book explores examples, terms, and programming needs relevant to those in the behavioral sciences, and helps readers perform virtually any computational function in solving their research problems. Principles are illustrated with usable code. Each chapter opens with a list of objectives followed by new commands required to accomplish those goals. The objectives also serve as a reference to help readers easily relocate a section of interest. Sample code and output and chapter problems demonstrate how to write a program and explore a model so readers can see the results using different equations and values. A website provides solutions to selected problems as well as the book's program code output and examples so readers can modify them as needed. The outputs on the website have color, motion, and sound.

Highlights of the new edition follow:

- Updated to reflect changes in the most recent version of MATLAB, including special tricks and new functions.
- More information on debugging and common errors as well as more basic problems in the rudiments of MATLAB to help novices get up and running more quickly.
- A new chapter on Psychtoolbox, a suite of programs specifically geared to behavioral science research.
- A new chapter on Graphical User Interfaces (GUIs) for user-friendly communication.
- Increased emphasis on pre-allocation of memory, recursion, handles, and matrix algebra operators.

Intended as a primary text for MATLAB courses for advanced undergraduate and/or graduate students in experimental and cognitive psychology and/or neuroscience, as well as a supplementary text for labs in data (statistical) analysis, research methods, and

computational modeling (programming), the book also appeals to individual researchers in these disciplines who wish to get up and running in MATLAB.

David A. Rosenbaum is a Distinguished Professor of Psychology at Pennsylvania State University.

Jonathan Vaughan is the James L. Ferguson Professor of Psychology and Neuroscience at Hamilton College.

Brad Wyble is Assistant Professor of Psychology at Pennsylvania State University.

MATLAB

for Behavioral Scientists

Second Edition

**David A. Rosenbaum,
Jonathan Vaughan, and Brad Wyble**

MATLAB® is a trademark of The MathWorks, Inc. and is used with permission. The MathWorks does not warrant the accuracy of the text or exercises in this book. This book's use or discussion of MATLAB® software or related products does not constitute endorsement or sponsorship by The MathWorks of a particular pedagogical approach or particular use of the MATLAB® software.

First published 2015
by Routledge
711 Third Avenue, New York, NY 10017

Simultaneously published in the UK
by Routledge
27 Church Road, Hove, East Sussex BN3 2FA

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2015 Taylor & Francis

The right of David A. Rosenbaum, Jonathan Vaughan, and Brad Wyble to be identified as authors of this work has been asserted by them in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reprinted or reproduced or utilised in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

Trademark notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data

Rosenbaum, David A.

MATLAB for behavioral scientists / authored by David A. Rosenbaum, Jonathan Vaughan, and Brad Wyble. — Second edition

pages cm

1. Psychology—Data processing. 2. MATLAB. I. Vaughan, Jonathan (Professor) II. Wyble, Brad.

BF39.5.R67 2014

150.285'536—dc23 2014003997

ISBN 978-0-415-53591-5 (hbk)

ISBN 978-0-415-53594-6 (pbk)

ISBN 978-0-203-11210-6 (ebk)

Typeset in Times and Courier New
by Apex CoVantage, LLC

Dedication Code

```
% Dedication.m
clc
for author = {'Brad' 'Jon' 'David'}
    authorstring = char(author);
    switch authorstring
        case 'Brad'
            Dedication.to = 'Elizabeth Spillman-Wyble';
            Dedication.features = ...
                {'inspiration','storytelling',...
                 'mastery of folklore',...
                 'extraordinary cooking'};
        case 'Jon'
            Dedication.to = 'Virginia Vaughan';
            Dedication.features = ...
                {'intelligence','strength of character',...
                 'unfailing support','generosity'};
        case 'David'
            Dedication.to = 'Judith Kroll';
            Dedication.features = ...
                {'brilliance', 'bravery', 'beauty'};
    end
    fprintf('%s dedicates this work to %s',...
            authorstring,Dedication.to);
    fprintf(' in grateful recognition of her ');
    for featurecount = 1:length(Dedication.features)-1
        fprintf('%s, ',...
                Dedication.features{featurecount});
    end
    fprintf('and %s.\n\n',Dedication.features{end})
end
commandwindow
```

Dedication Output

Brad dedicates this work to Elizabeth Spillman-Wyble in grateful recognition of her inspiration, storytelling, mastery of folklore, and extraordinary cooking.

Jon dedicates this work to Virginia Vaughan in grateful recognition of her intelligence, strength of character, unfailing support, and generosity.

David dedicates this work to Judith Kroll in grateful recognition of her brilliance, bravery, and beauty.

This page intentionally left blank

Contents

Preface	ix
Acknowledgements	xii
About the Authors	xiii
1 Introduction	1
2 Interacting With MATLAB	22
3 Matrices	38
4 Calculations	62
5 Contingencies	99
6 Input-Output	128
7 Data Types	156
8 Modules and Functions	182
9 Plots	208
10 Lines, Shapes, and Images	248
11 Animation and Sound	287
12 Enhanced User Interaction	304
13 Psychtoolbox	323
14 Debugging	355
15 Going On	370

References	375
Commands Index	377
Name Index	381
Subject Index	382

Preface

The first edition of MATLAB for Behavioral Scientists (published in 2007) was the result of a rebellious thought. The prevailing view before then was that most behavioral scientists shouldn't or couldn't write their own computer programs. This irked the first author, who decided to pursue the notion that all behavioral scientists, including students in the relevant fields (psychology, cognitive and affective neuroscience, economics, and so on), could and should learn to program for themselves.

Behavioral scientists need to be able to program as much as scientists in other fields. They need to be able to program to do whatever they want, computationally speaking, without having to rely on the kindness of strangers or the largesse of granting agencies to pay others to program for them.

To give some examples, a behavioral scientist—a behavioral economist, say—wishing to model decision making should be able to roll up her sleeves and graph data showing observed and expected data in the way she prefers. A personality psychologist interested in designing a new questionnaire requiring a special computer interface should be able to pursue that aim. A psychotherapist wanting to model changing relations between members of a family should be able to characterize that process with custom-made animations that show network links with dynamically changing thicknesses and colors, growing and shrinking over time, if that's what she wants. A cognitive psychologist interested in setting up and conducting behavioral experiments should be able to create any kind of stimuli and response recording capabilities he or she cares to, not being limited by what's possible with off-the-shelf commercial products.

This book is meant to help behavioral scientists (and especially students entering this field) to do these things. The authors of this book assume you have no prior familiarity with computer programming, and we assume you have no knowledge of mathematics beyond what is generally learned in high school. The text is meant to be as friendly and encouraging as possible. Our aim is to draw you in and help you feel comfortable within a domain that may at first seem foreign and maybe even scary.

Programming can be humbling. If you set out to learn to program, you should prepare yourself emotionally as well as intellectually for what will happen because you will be dealing with an unfeeling machine. It takes a tough hide to believe you have a program that does what you want, only to discover that the program doesn't run, generates unexpected results, or produces outputs that seem reasonable at first but then turn out to be wrong. Everyone who has programmed has gone through this, including the authors of this book, so don't feel like you need to be able to program perfectly. No one does!

Programming needn't be unpleasant, however. The attitude to have is to keep an open mind about the value of mistakes. If you treat errors as windows for improvement, you will learn a lot. Availing yourself of that learning, when you see a program work and especially

when it does something that, to your knowledge, has not been done before, can let you feel rightly proud of your achievement.

There are many computer programming languages. Why is this book about MATLAB? MATLAB (short for Matrix Laboratory), is a commercial product of a company called The MathWorks (Natick, Massachusetts), for which we authors do not work and have no commercial connection. The following, therefore, can be taken as our honest opinion of their product: MATLAB is a simple yet powerful language for computer programming. It has an active community of users, engaged in many branches of science and engineering. One of MATLAB's most attractive features is that it offers high-level commands for performing calculations with large as well as small data sets and for generating publication-quality graphics. Another attraction of MATLAB is that it allows for the presentation of stimuli and the collection of responses with precise timing. Yet another attraction is that MATLAB is platform-independent. It runs on PCs, Macs, and Linux machines. For these and other reasons, MATLAB is a very good language for behavioral scientists. A growing number of behavioral scientists, along with neuroscientists, engineers, and investigators in other disciplines, have therefore chosen to learn MATLAB. Owing to the health and vitality of the MATLAB programming community, it is likely that more and more people will want to learn MATLAB in the future. You will be part of that active community if you choose to dive into the material provided here.

How did it come to pass that there is a second edition of this book? As is always true of a second edition, its predecessor was successful enough to keep the work alive, but changes in the field suggested a face-lift was needed. Among the needed changes was the appearance of other MATLAB books for psychologists and neuroscientists (Fine & Boynton, 2013; Madan, 2014; Wallisch et al., 2009), which are welcome additions, though they are different in style, tone, level of coverage, and organization from the first edition of this book (but not so perfect, in our view, that they obviate this second edition).

As the author of the first edition (Rosenbaum, 2007) contemplated the second edition, he realized that the process of revising and updating the book would benefit from the involvement of his long-time friend and collaborator, Jonathan Vaughan, the James L. Ferguson Professor of Psychology and Neuroscience at Hamilton College. Jon has decades of experience with computer programming. He has served as the editor of *Behavior Research Methods, Instruments, & Computers*, a peer-reviewed publication of the Psychonomic Society. The first author basically learned MATLAB from Jon. He continued to learn from Jon in preparing this second edition.

When Jon agreed to join in, he and David began to map out the ways the second edition would differ from the first. Among the things they agreed to were the following: (1) All known errors in the first edition would be corrected; (2) more would be said about debugging; (3) more problems would be given, including problems that would help students confront very basic issues in the rudiments of MATLAB; (4) solutions to selected problems would appear with downloadable code on the book's new website (www.routledge.com/9780415535946) rather than in the back of the book to allow for more extensive code, updating of the programs if necessary, and addition of new programs as needs and curiosities arose; (5) there would be a tutorial on designing Graphical User Interfaces, or GUIs, which enable a user to interact with a program using graphics to run experiments within MATLAB; (6) there would be a tutorial in designing experiments using Psychtoolbox, a freely available MATLAB toolbox that is specifically geared to behavioral science

research; and (7) special tricks and new functions, developed or discovered since 2007, would be featured, including several developed by the authors to solve sometimes thorny problems that arise in data collection and data presentation.

In preparing the second edition, Jon and David made these changes while retaining the main organization of the book's first edition. As before, readers are ushered to the material slowly and in as a welcoming a way as possible, with more specialized topics coming as the chapters continue. Also as in the first edition, there is continued use of a style that worked well before—introducing a new problem or challenge, presenting associated code, and then presenting the output. In addition, as in the first edition, each chapter starts with a list of things to be done followed by commands that get them done. These start-of-chapter lists let you use the book as a reference once you understand the basics of MATLAB. Thus, after you have worked your way through the book, you will be able to turn to a section and quickly get the detailed information you need to complete the programming task you are undertaking. All the commands are listed as well in a single Command Index near the back of the book, another innovation of the second edition relative to the first.

Another way we have made the text as user-friendly as possible is to update the website for this book. On this site, you will be able to find and copy the programs and program outputs in this volume. The outputs on the website have color, motion, and sound, whereas those modalities are absent from the printed edition.

As shown in the list of new features, the second edition has a chapter on Psychtoolbox. This is a free, popular, MATLAB-based toolbox for running behavioral experiments. Neither Jon nor David had used Psychtoolbox before, simply because it wasn't essential for their work. It happened, however, that Brad Wyble, a newly hired faculty member in the Penn State Psychology Department (the department where David works), had extensive experience with Psychtoolbox. Jon and David invited Brad to prepare a chapter for the book on Psychtoolbox, and, to their great satisfaction, he agreed.

Brad's area of expertise is vision, the domain of behavioral science in which, it happens, Psychtoolbox is used the most. With his extensive background in computer science—Brad was a computer science major as an undergrad and did research in computer science labs after completing his PhD at Harvard—he proved to be a wonderful addition to the team. His involvement in the book was limited to the one chapter he prepared, plus his review of this Preface, as per the agreement he made with Jon and David. Any errors in the book, then, outside of the Psychtoolbox chapter and the Preface are not due to Brad. By the same token, any errors in the Psychtoolbox chapter and in the Preface are as much Jon's and David's fault as they are, or might be, Brad's. In general, any mistakes rest squarely with Jon and David, or most especially David, who, after having had several years to mull over the transition from the first edition to the second, should have gotten things right by now!

The last thing we want to say in this preface echoes what we say in the main text about responsiveness to feedback. It is fine to be open to feedback from a *computer*, as we urge you to be, but it is also good to be open to feedback from *people*. If you spot something that you think could be better, please let us know. If you have suggestions for things to include in a future edition, give us those suggestions. If you want help with your programming, we cannot serve as consultants to you. We appreciate understanding on that last point. To get in touch with us, you can use one or more of our e-mail addresses: dar12@psu.edu, jaughan@hamilton.edu, or bpw10@psu.edu. We hope you will find this book useful.

Acknowledgements

There are others who deserve praise and thanks for their contributions, direct and indirect. First, we express our appreciation to the students who took the MATLAB courses offered by David at Penn State and by Jon at Hamilton, and who also were exposed to MATLAB by Brad. Teaching these students helped us see which programming concepts are transparent and which are less so.

Several students in our classes and in our labs played especially important roles in helping us hone our MATLAB instruction. We thank Penn State students Max Bay, Katie Chapman, Chase Coelho, Rajal Cohen, Samantha Debes, Jeff Eder, Jason Gullifer, Lanyun Gong, Derek Henig, Joe Santamaria, Garrett Swan, Matt Walsh, and Robrecht van der Wel. We thank Hamilton College students Deborah Barany, Julia Brandt, Hallie Brown, Drew Linsley, Ramya Ramnath, and Anthony Sali. Others who provided valuable feedback are Debra Boutin, Gillian Dale, Mike Frederick, Michael Romano, and Doug Weldon. Mario Kleiner provided helpful information about Psychtoolbox.

We also wish to thank the reviewers who provided feedback on the revision plan: Simon Farrell, University of Bristol, UK; Alen Hajnal, University of Southern Mississippi, USA; and an anonymous reviewer.

This book was completed while the first author was on sabbatical in Los Angeles, at UCLA and USC, where he was supported in part by a fellowship from the John Simon Guggenheim Memorial Foundation. Brad Wyble's research was supported at the time of this writing by NSF grant BCS #1331073. Jon's research has been supported by grants from the National Science Foundation and the National Institutes of Health, as has David's. We all appreciate this support, not to mention the support of the institutions that have paid our salaries.

We also wish to express our thanks to Paul Dukes at Psychology Press/Taylor & Francis, who was instrumental in opening the door for the second edition of the book. Paul called on his colleague, Debra Riegert, to work with us to bring the work to completion. Debra was responsive and helpful at every stage. We appreciate her help as well as the further assistance of Angela Halliday at Routledge/Taylor & Francis, who helped with the book's and website's production.

About the Authors

David A. Rosenbaum is a cognitive psychologist whose main interests are human perception and performance. His main research contribution has been joining cognitive psychology and motor control. Rosenbaum attended public schools in Philadelphia and then attended Swarthmore College (B.A., 1970–1973) and Stanford University (Ph.D., 1973–1977). He worked at Bell Laboratories (1977–1981), Hampshire College (1981–1987), and the University of Massachusetts, Amherst (1987–1994). He has been at Pennsylvania State University since 1994, where he was named Distinguished Professor of Psychology in 2000. Rosenbaum was a recipient of a National Science Foundation Graduate Fellowship (1973–1976), a National Institutes of Health Research Career Development Award (1985–1990), and a National Institutes of Health Research Scientist Development Award (1992–1997). His work has been supported by grants from the National Science Foundation (NSF) and the National Institutes of Health, as well as grants from the Dutch, French, and German equivalents of NSF. Rosenbaum is a Fellow of the American Association for the Advancement of Science, the American Psychological Association, the American Psychological Society, and the Society of Experimental Psychologists. He served as Editor of *Journal of Experimental Psychology: Human Perception and Performance* (a publication of the American Psychological Association) from 2000 to 2005. He was awarded a Guggenheim Foundation Fellowship in 2012 for the 2013–2014 academic year. Besides being the author of the first edition of this book, David is the author of a textbook on motor control [Rosenbaum, 2010] and the author of a book applying Darwin’s theory of natural selection to cognitive psychology [Rosenbaum, 2013].

Jonathan Vaughan is a broadly trained experimental psychologist (B.A., Swarthmore College, 1962–1966; Ph.D., Brown University, 1966–1970) whose research interests focus on the planning and execution of motor actions, eye movements and attentional processes, human and animal learning, and cognitive neuropsychology. He has taught at Hamilton College since 1971. His work with David Rosenbaum and Ruud G.J. Meulenbroek, initiated under an AREA grant from the NINDS, has produced computational models of reaching, grasping, tapping, and manual circumvention of obstacles. Other research support has come from the NSF and NIMH. Vaughan has published more than 60 journal articles and book chapters, and given more than 100 research presentations, many in collaboration with Hamilton undergraduates. He has contributed in many ways to computer applications in psychological research, including tutorial materials for the use of PsyScope and SPSS. He edited the Psychonomic Society’s international quarterly, *Behavior Research Methods, Instruments, and Computers* [1994–2004] and founded the Psychonomic Society’s Archive of Norms, Stimuli, and Data, an online repository of computer programs, data, and stimulus norms that has served as an important resource for researchers in the field.

Brad Wyble studies attention, perception, and memory. He attended public schools in Lancaster, Pennsylvania, after which he obtained a B.A. in computer science from Brandeis University (1991–1995) and a Ph.D. in psychology from Harvard University (1996–2003). He was a postdoctoral fellow at the University of Kent in Canterbury, England (2003–2007), University College, London (2007), and MIT (2007–2009). He was subsequently an assistant professor at Syracuse University (2009–2012) and is now an assistant professor at Pennsylvania State University in the Department of Psychology. Wyble was a recipient of a National Science Foundation Graduate Fellowship (1997–2000), he was a Sackler Fellow (2001–2002), and he has been supported by grants from the National Science Foundation, the Office of Naval Research, and the National Institutes of Health. He serves as a consulting editor for the *Journal of Experimental Psychology: Human Perception and Performance*, and as an associate editor for the journal *Frontiers in Cognition*.

1. Introduction

This chapter covers the following topics:

- 1.1 Getting oriented
- 1.2 Getting an overview of this book
- 1.3 Understanding computer architecture
- 1.4 Programming principles
- 1.5 Deciding if a program is needed and whether you should write it
- 1.6 Being as clear as possible about what your program should do
- 1.7 Working incrementally
- 1.8 Being open to negative feedback
- 1.9 Programming with a friend
- 1.10 Writing concise programs
- 1.11 Writing clear programs
- 1.12 Writing correct programs
- 1.13 Understanding how the chapters of this book are organized
- 1.14 Using the website associated with this book
- 1.15 Obtaining and installing MATLAB
- 1.16 Acknowledging limits

1.1 Getting Oriented

Computers are vital in every branch of science today, and behavioral science is no exception. When behavioral scientists use computers to obtain responses in questionnaires, present visual stimuli, display brain images, generate data graphs, or write manuscripts, their ability to make efficient progress in their research depends largely on their ability to use computers effectively.

Many specialized computer packages let behavioral scientists do their work, and each one takes some time to learn. It is useful to know how to use these specialized packages, but it is also tantalizing to consider the possibility of learning how to program for yourself. The reason is that all specialized computer packages rely on underlying code, and knowing how to generate such code yourself can allow you to be self-sufficient or nearly so in your own research.

Suppose, for example, that you want to develop a mathematical model of some cognitive process. It is convenient to be able to write a program that lets you explore the mathematical model freely, seeing the results obtained with different equations, different parameter values, and so on. Similarly, to analyze data in ways that would be cumbersome with existing spreadsheet applications, it is refreshing to be able to write the analysis program to your own specifications. For example, to view graphs of obtained or theoretical data in a variety of forms, it is useful to be able to generate the graphs quickly and easily, however you please, not just as stipulated by an existing graphics package.

The computer language introduced here, MATLAB, provides you with these capabilities. MATLAB is available from The MathWorks (www.mathworks.com), a company with which

we authors have no affiliation. MATLAB has become popular in several branches of engineering and science, including behavioral science. Nonetheless, to the best of our knowledge, no book has appeared about MATLAB that is written specifically with behavioral scientists in mind. Nor for that matter has a book come out for behavioral scientists about any other general-purpose programming language. The need for such a volume motivated the first edition of this book. Its positive reception encouraged us to revise the text and expand the coverage in this second edition.

Will it be worth your time to read this book? Once you have gone through the text and generated your own MATLAB programs based on the material presented here, you should have enough programming skill to do most of what you need to for your own behavioral research needs. Most importantly, a working knowledge of MATLAB will allow you to perform some analyses that would be tedious, difficult, or impossible otherwise. In addition, you will be able to understand and build upon the work of colleagues who use MATLAB in their work.

You will probably find this book most useful if you use it in two stages. In the first, you will want to go through it, or the parts of it most relevant to your needs, in considerable detail, working problems and developing the hands-on skills that will make you a MATLAB *user*, not just a MATLAB *appreciator*. In the second stage, you will be able to rely on the book as a reference, turning quickly to those sections that provide examples you can adapt for your own programming needs.

To make the book as useful as possible as a reference source, we have designed it so you can get the examples you need quickly and easily. You can do so by turning to the opening page of any chapter and finding there a list of things you may want to do. Beneath that list is a compendium of associated commands. The text itself provides examples you can adapt for your own purposes. You can copy those examples by hand into your own programs, or, to avoid typographical errors, you can copy and paste them from the website associated with this book (www.routledge.com/9780415535946), where the programs and their outputs are available, along with the solution to selected problems. Finally, the list of commands introduced in each chapter is listed as well in the Commands Index.

1.2 Getting an Overview of This Book

Acquiring a new skill such as computer programming can be daunting, so it helps to have an overview of what you can expect as you proceed. Here, then, is a roadmap of the contents of this book. Besides signposts, we also provide brief explanations of the goals of each chapter.

1. **Introduction.** By reading the present chapter, you will learn more than you may already know about how computers work and what computer programming languages do. You will also learn about the ways you should approach computer programming. Finally, by reading this chapter, you will understand how this book is organized. That information can help you use the book efficiently.
2. **Interacting With MATLAB.** By delving into the second chapter, you will learn how to activate MATLAB's windows in order to open, edit, save, and run MATLAB programs.

3. **Matrices.** By studying the third chapter, you will learn how MATLAB enables you to store and access data. Briefly, MATLAB lets you store data in matrices consisting of one or more rows and one or more columns. Matrices are so fundamental to MATLAB that the name of the language is actually short for “Matrix Laboratory.” You can think of a two-dimensional matrix (one having both rows and columns) as analogous to the rows and columns in a spreadsheet.
4. **Calculations.** Computers are good at calculating. Chapter 4 shows how to get your computer to carry out calculations with MATLAB.
5. **Contingencies.** One of the main purposes of a computer program is to perform different actions depending on existing conditions. The logic of a program involves not only calculations but also decision making, such as evaluating variables differently (or not evaluating them at all), depending on their values.
6. **Input-Output.** Chapter 6 shows you how to control your computer’s interactions with the external world. By studying Chapter 6, you will be able to design programs that let you create dialogs with users, including participants in behavioral studies, and to read and write data to and from external files.
7. **Data Types.** One of the biggest challenges in using computers in research is determining how best to represent the data you are working with. It is important to understand what data types are available in MATLAB so you can choose and manipulate your data types accordingly.
8. **Modules and Functions.** Simple programs are usually easy to understand, but when they become more complex it often helps to deal with them in chunks. Some higher level structure is often helpful. Chapter 8 shows you how to write programs that have this property. Those programs often have stand-alone modules and functions. Such modules and functions can be called by a variety of programs. Using modules and functions can help you approach programming from a top-down rather than a bottom-up perspective. Modules and functions can also facilitate the reuse of programs in the future.
9. **Plots.** The ability to generate and manipulate complex graphics for the exploration and presentation of data is widely regarded as one of the special strengths of MATLAB. Chapter 9 exposes you to those strengths by showing you how to make line graphs, bar graphs, and other types of graphs that are suitable for professional presentations and publications.
10. **Lines, Shapes, and Images.** Here you will learn how to create, import, or reshape lines, shapes, and other images that can either stand alone or be included in graphs. Chapter 10 will also show you how to generate three-dimensional graphs.
11. **Animation and Sound.** Chapter 11 builds on the static graphics of the tenth chapter to manipulate figures using simple animation techniques, generate movies, and generate auditory stimuli.
12. **Enhanced User Interaction.** When you think of a typical computer application, what comes to mind is how the program interacts with the user, typically through graphics, the keyboard, the mouse, or touchscreen. Chapter 12 introduces you to some of the tools available in MATLAB for user interactions.

13. **Psychtoolbox.** For real-time work, there are some features that MATLAB ordinarily lacks that are needed for precise and flexible stimulus presentation and data acquisition. Chapter 13 describes a sophisticated extension to MATLAB, *Psychtoolbox*, which adds features to facilitate research using MATLAB, especially in vision research. This chapter also touches on related packages of interest to behavioral scientists in related areas.
14. **Debugging.** Programs often have bugs because, for better or worse, programming is often a trial-and-error process. While it is hard to know in advance how to address every possible bug, it is possible, based on the authors' many goofs of their own, to convey advice about debugging techniques which you may find useful. These are offered in Chapter 13 . . . oops, Chapter 14 (☺).
15. **Going On.** Chapter 15, the last chapter of the book, provides pointers for going further with MATLAB. This chapter also directs you to other resources you may want to draw on.

A lot of material will be covered in this book. Do you need to go through all of it? If you have no need to play sounds, show animations, or generate three-dimensional graphics, you may safely ignore large parts of Chapters 9 through 13, though leafing through these chapters may help you overcome any prejudices or fears you might have regarding these topics. At the same time, there are chapters you cannot avoid, at least if you don't want to emerge from this book the way Woody Allen emerged from his speed-reading of Tolstoy's epic novel, *War and Peace*. "It was about Russia" was all he could recall.

The truly essential chapters of this book are Chapters 2 through 5. You cannot go on to the later chapters and expect to have control of your programs if you don't have command of the material in Chapters 2 through 5, and the only way to gain that command is to work your way through the examples and exercises slowly and carefully. We promise that even if you think you understand how things work, the only way to be sure is to try them out and expose yourself to the feedback you will receive.

As you gain expertise, Chapters 6 through 8 will allow you to write more sophisticated code. Chapters 9 through 13 will provide you with specialized tools for your work and enjoyment. And Chapter 14, as already mentioned, will suggest ways to help you debug efficiently.

A word of advice: Don't hesitate to revisit earlier sections of the book as you move through it. No one remembers perfectly, and no one understands material quite as fully the first time as in revisits. Your understanding of what may seem very obscure the first time through will be enhanced by the top-down knowledge and context you will acquire touring later material.

1.3 Understanding Computer Architecture

As a first step toward learning to program, it can be helpful to know a bit about computer architecture. Knowing about the main components of a computer can help you understand what features of the environment your program must deal with.

All working computers have five basic elements. As shown in Figure 1.3.1, these are (1) input devices (not only the conventional keyboards and mice, but also the microphones, response buttons, and video and voltage recorders that are useful in the laboratory); (2) output devices (screens, printers, loudspeakers, etc.); (3) storage devices (hard disks, thumb drives, DVDs, the “Cloud,” etc.); (4) primary memory; and (5) the central processing unit. The first three components should need no further explanation. The last two components merit more discussion.




Figure 1.3.1

Primary memory (item 4 on the list) is like human or animal working memory. Its contents are currently active information. The amount of information that can be kept in this active state is limited, both in biological agents (humans and animals) and in computers. The amount of information a computer can maintain in primary memory is hardware dependent.

Because the capacity of primary memory is limited, it is important to be mindful of the amount of information a computer can keep active at once. The amount of information made active by a program, such as one written in MATLAB, depends on the number of variables that are declared and the number of bits (the number of 1s and 0s) required to represent each variable.

Essentially, there are three ways of using primary memory efficiently: (1) defining just the variables that are needed; (2) clearing variables once they are no longer needed; and (3) defining the types of the variables so the amount of memory initially reserved for them is large enough but not substantially larger than needed. We will return to these topics in Chapter 7 (“Data Types”).

Returning to the components of computer architecture, the fifth component is the central processing unit. This is the part of the computer that executes instructions. For present purposes, the central processing unit, or CPU, can be likened to consciousness, for which, it is said, only one thought can exist at a time (James, 1890). The same can be said of a computer’s CPU. It can handle only one instruction at a time, at least in a conventional digital computer. Handling just one instruction at a time is called *serial* processing. Handling more than one instruction at a time is called *parallel* processing.

Serial processing can occur at high rates in modern computers. For example, the computer on which this text was prepared (a Dell laptop) runs at 2 gigahertz (2 billion cycles per second).

Regardless of the speed at which a CPU runs, serial processing imposes constraints on the kinds of programs you can run, and therefore write, in MATLAB. Suppose, for example, that you want to find the largest value among a set of numbers. Parallel processing is a natural way to solve this problem. If the values are plotted as in Figure 1.3.2, for example, a brief glance at the bars lets you pick the biggest one. The tallest bar seems to jump out at you. Once it does, you can look down to find the associated element (element 3 in this case), or you can look to the left to find the largest value (39 in this case).




Figure 1.3.2

You might object that parallel evaluation of the heights of all the bars in this case is not actually possible, and even it were for this particular figure, it wouldn't be for all other sets of numbers, such as those whose largest values are similar. You might also say that the method outlined above is not a truly parallel process because distinct stages are associated with looking down the tallest bar or looking sidewise from the top of the tallest bar. These objections are well taken, especially considering that serial processing is inescapable in MATLAB, at least in a program that uses MATLAB in its usual configuration. To sort values or do anything else in MATLAB, everything must be done one step at a time (serially). Knowing this can help you approach the task of programming. (Many recent computers have multiple processors, or *cores*, that make parallel computing possible. Advanced users can take advantage of these to speed complex computations by having two or more cores compute different things at once, using additional tools available from The MathWorks. If you are beginning your programming skills with this book, you can safely save parallel programming for another time.)

1.4 Programming Principles

How should you approach the task of programming? We have come to believe in the following principles:

- Decide if a program is actually needed and, if so, whether you should write it.
- Be as clear as possible about what your program should do.

- Work incrementally.
- Be open to negative feedback.
- Program with a friend.
- Write concise programs.
- Write clear programs.
- Write correct programs.

Consider each of these principles in turn.

1.5 Deciding If a Program Is Needed and Whether You Should Write It

The first principle is less obvious than you might suppose. Consider the problem discussed above (finding the largest of a set of values). The numbers corresponding to the bars in Figure 1.3.2 are as follows:

7 33 39 26 8 18 15 4 0

Do you need a computer program to find the largest of these values? Obviously not. You know that the largest of these numbers is 39 and that this largest number occupies the third slot in the series. If you only had to find the largest value in this particular array, you would be foolish to write a program for this task, except as an exercise. On the other hand, if you were quite sure you would often need to find the largest number in each of a large number of arrays of unpredictable sizes, writing a program would make more sense. A program is useful, then, for performing a well-defined task that would be too taxing to perform by hand.

The second part of the first principle, whether you should write the program yourself, also deserves comment. If you decide you need a program, it may or may not make sense for you to write the program yourself. Why should you write a program for a task if someone else has done so before?

Our answer to this question is analogous to the answer a math teacher might give to a rebellious student: “Why should I prove this theorem if it’s been proved before?” “Practice makes perfect,” the teacher may reply. He or she may go on: “Even if true perfection is beyond your reach, practice will increase the chance of your proving something new yourself.”

Our view of programming is the same. You might be able to locate programs that already do things you need to, and it may make sense for you to use those programs, especially for problems that seem very complicated or that are beyond your technical ability. But the more practice you get programming, the more likely it will be that you will be able to generate programs that either solve new problems or solve old problems in new ways. Don’t be discouraged if it takes an hour or more to get your first “real” program up and running, even if you might have done the same computation by hand in a minute or less. As you develop

programming expertise, you will become more efficient and productive, and you'll be able to apply your new skills to other problems.

1.6 Being as Clear as Possible About What Your Program Should Do

If you decide that you need a program and that you should write it yourself, you will need to be as clear as possible about what your program should do. This is easier said than done. Thinking through the workings of a program can be one of the hardest aspects of programming, even harder in some cases than getting the syntax right.

Return to the problem of finding the largest value in an array. It turns out that MATLAB provides a program (or more precisely, a *function*), called `max`, that lets you find the maximum of a set of values (see Chapter 4). You can use this function to get the largest value in a matrix without having to reinvent the function yourself. Nevertheless, it is worth thinking through the way you would identify the largest value in an array. Working through this example—however simple it may seem—will help you begin to “think programmatically.”

To think through what a program must do to find the largest value in an array of numbers, imagine that you have a row of numbers like the one above, but you can only see one of the numbers at a time—say, by sliding the hole in a card across the row. Under this circumstance, you can determine the largest value by finding the largest value *so far*. If you were actually doing this, you'd first place the hole in the card over the first number, which is 7. Then, you'd remember that 7 is the largest value you've seen, and move the card to reveal the 33. Thirty-three is larger than 7, so now you'd note that 33 is the largest number you've seen, and you'd move the card again. After seeing 39, you would revise the largest number seen to that value. Continuing and not encountering any number larger than 39 for the rest of the series, that would be the number you'd report.

Now translate this algorithm into a program. Assign some very small value to a variable called, for instance, `Largest_Value_So_Far`. Then, proceeding from left to right, every time you encounter a value larger than `Largest_Value_So_Far`, reassign that new value to `Largest_Value_So_Far`. After you have evaluated the last item on the list, `Largest_Value_So_Far` will be the largest of all the values.

Here is a flow chart for the procedure, along with some other items you'd need to get the job done. One of these other items is telling the program how many values there are in the list. We give the list the name `V`. There are $n = 9$ values in `V`.

Another thing that needs to be done is initializing `Largest_Value_So_Far` to an extremely small value, namely, minus infinity (which can be expressed in MATLAB as `-inf`). We do this because whenever a new number is tested, it must be compared to some prior value. Starting with `-inf` ensures that the first value will be called the largest provided it is larger than `-inf`. It may stay that way if no larger value comes along.

The third thing that needs to be done is providing an index, `i`, for each successively encountered value in `V`. An index for a value is the position of the value in the matrix. For the first item, $i = 1$, for the second item, $i = 2$, and so forth. Initially, `i` is set to 0. Each time a new number is compared to `Largest_Value_So_Far`, the variable `i` is incremented by 1, until `i` is greater than `n`. The i -th value of `V`, denoted `V(i)`, is assigned to `Largest_Value_So_Far` if `V(i)` is larger than the current value of

`Largest_Value_So_Far`. When i is larger than n , the program stops and the value of `Largest_Value_So_Far` is printed out.




Figure 1.6.1

A flowchart like this can serve as the conceptual foundation for the code needed to get a computer to find the largest value in an array. You don't *have* to draw a flowchart before you write MATLAB code, however. Some people only imagine flowcharts or the steps corresponding to them. Drawing flowcharts in your head obviously gets easier as you get more practice with programming. Early in practice, however, it is advisable to sketch the steps your programs will follow.

How do you come up with a flowchart or its corresponding steps in the first place? The honest answer is that no one knows. Anyone who could give the answer would, in effect, know how thoughts originate, and no one at this time has a clue about that. If you solve this problem, a Nobel Prize awaits you.

You can, however, consider some practical advice about how to come up with the procedures for computer programs. One suggestion is to talk out loud as you imagine yourself doing the task you wish to program, step by step, much as we did with the imaginary card above. Talking out loud may enable you to make explicit whatever implicit knowledge you bring to bear as you do the task, as if you were explaining the task to a friend. Hearing your own words will also help you identify those things you're not clear about. If you hear yourself say, "OK, next I'll somehow figure out which of the values might be OK based on some criterion I can't quite articulate but I have a vague feeling about," then you're not quite ready to write all the code you need. Ultimately, you'll need to be completely explicit about the instructions your programs contain. Relying on a miracle just won't work, and the reason, just to be explicit, is that computers, for all their speed, are ignorant and inflexible. They do exactly and only what they're told to do.

This is one way in which programming is very different from other forms of communication. When you speak to other people, you assume—usually correctly—that they have some knowledge that lets them fill in missing information. Not so with computers, or at least conventional computers given stand-alone programs. Writing successful computer programs requires a degree of explicitness that is unparalleled in other aspects of human experience. This is one reason why learning to write computer programs can be challenging. On the other hand, being explicit to the point that a computer can carry out instructions may sometimes carry over well to other things you do, like writing papers or reaching agreements with others about who will do what in connection with some project.

1.7 Working Incrementally

Another challenge of programming is translating your procedural ideas into language the computer can understand. Here it is useful to work incrementally. By this we mean you should build your program a little at a time, making sure each part works before you go on to another part that depends on what you’ve just written. You should build your program the way a reliable contractor builds a house, by making sure the foundation is solid before the basement is added, by making sure the basement is solid before the first floor is added, and so on. During program development, you will often find it useful to generate intermediate output to verify that each step works as expected. You may later inhibit that output when the program is completed and is no longer needed. Think of this incremental programming process as the digital equivalent of the ancient woodworking adage (attributed to John Florio, 1591), *Alwaies measure manie, before you cut anie* (“Measure twice, cut once.”).

When you’re reasonably sure your program works, and before you add another component or make other significant changes, save the program with a file name unique to the last working version. The moment you prepare to make changes to the program, save the file with a new name or version number before putting in any changes. Follow the American folk adage, “If it ain’t broke, don’t fix it.” Too often, attempts to further develop a program *will*, in fact, break it, or otherwise reveal some weakness in it, and you might want to go back to an earlier version. You’ll be glad you have one!

Remember, too, that computer storage is cheap. There is no harm in having a folder full of documents called `Max_Program_01.m`, `Max_Program_02.m`, `Max_Program_03.m`, and so on. It may be that the version you’ll use for actual work is `Max_Program_101.m`. There is nothing wrong with such a high number. You can tuck away the earlier versions in a sub-folder until you’re sure you’ll never need to look back. Having sequential versions of a program in development makes it easy to compare the changes. In this connection, it is useful to note that MATLAB has a comparison tool that highlights all differences between two versions of a program, similar to “track changes” in Microsoft Word.

1.8 Being Open to Negative Feedback

How can you tell if your program works? As you consider this question, one attitude should rule over all others: *Be open to negative feedback*. If you treat negative feedback as a help rather than a hindrance, you will become a better, and certainly happier, programmer than if you treat negative feedback in a negative way.

The research of psychologists Carol Dweck and Janine Bempechat (1980) is relevant in this regard. Dweck and Bempechat distinguished between people who take negative feedback

as signs of their lack of talent (*entity* learners) and people who treat negative feedback as cues for ways to improve their performance (*incremental* learners). It is important while programming to have the attitude of an incremental learner rather than an entity learner. You will learn more if you take negative feedback constructively than if you read such feedback as a sign that you weren't "cut out" for programming. MATLAB will not give you an error message that says

```
??? You don't deserve oxygen!
```

A more likely message is something prosaic like

```
??? Subscript indices must either be real positive integers
or logicals.
```

You might get an error message like the latter one in response to code such as

```
Reaction_Time_For_Trial(0) = 680;
```

All you have to do here is appreciate that it makes no sense to have the zero-th element of an array. An array can have a first element, a second element, a third element, and so on, but it can't have an element numbered zero. Whether the 0 was entered in the code based on a misunderstanding or simply as a typo, you can correct the error without indicting your genes. If when you typed 0, you were referring to the first trial, you can replace the 0 with a 1 and all will be fine:

```
Reaction_Time_For_Trial(1) = 680;
```

One reason for saying these things is that it bears remembering that the error messages you receive while programming come from a machine, not from a person who knows what you are trying to say. When you receive an error message, it will help you to take the message as a piece of advice. Over time, you will get fewer error messages concerning low-level aspects of coding (e.g., when you have an unequal number of opening and closing parentheses in a line of code), and you will learn what the error messages mean. More about error messages and debugging (correcting your programs) will come later in the text.

Over time you will also learn to guard against disaster when you program. We encourage you to do so by writing programs that are resilient rather than brittle. If you write a program that crashes or yields crazy results when it gets input of a different sort than what you anticipated, your program won't be of much good. For example, if you write a program that is used to collect questionnaire data, and a participant types in an age of -83, that could wreak havoc with subsequent data analyses. It doesn't matter why the participant put a minus sign in front of his or her age (if he or she is actually 83). Perhaps the participant thought this might help you see the number more clearly, perhaps it was just a typo, or perhaps the participant thought he or she was being cute. The point is that you must anticipate such eventualities. All sorts of things can go wrong when a program is being run. A good programmer guards against such eventualities. In this sense, being open to negative feedback means more than not letting your feelings be hurt when the computer beeps because you left out a punctuation mark or because you mistyped the name of a function. Responding constructively to negative feedback also means being open to all sorts of

unwanted events and building safeguards into your programs so you're not confronted with bogus results later on.

The final sense in which it is important to be open to negative results is that you should not be complacent when your program runs and gives you results, especially beautiful ones, that cause you to blush with quixotic pride. Here is an example.


The numbers 1 through 8 are assigned to a matrix called `x`. These numbers are session numbers, which comprise the independent variable of a fictional behavioral science study. The dependent variable is `y`, a set of fictional scores. After `x` and `y` have been defined, a command is used to `plot` the data. This command ends with a special instruction, in quotes, to plot the data in black (`k`), using circles (`o`), with connecting lines (`-`). Within the plot command, you accidentally (or on purpose for this example) tell MATLAB to plot `x` along the horizontal axis and to plot `x` along the vertical axis, rather than telling MATLAB to plot `x` along the horizontal axis and `y` along the vertical axis. Three more lines of code follow. One sets the limits of the `x` axis to ensure that the first point is plotted (a need that arises for this particular graph). The second specifies the label for the `x` axis, using the `xlabel` command. The third specifies the label for the `y` axis, using the `ylabel` command. (More details about these commands will be given in Chapter 9. You can just skim over them here.)

Code 1.8.1:

```
x = [1 2 3 4 5 6 7 8];
y = [0.39 0.47 0.60 0.21 0.57 0.36 0.64 0.32];
plot(x,x,'ko-')
xlim([0 9])
xlabel('Session')
ylabel('Score on Test')
```

When you look at the output, you are impressed by the beauty of the results.

Output 1.8.1:




Before calling a press conference, however, it would be advisable for you to check your work. In this case, the results look too good to be true, and in fact, they are. An error was made. Once the error has been found and fixed (with a comment inserted in the program accordingly), the results look quite different.

Code 1.8.2:

```
x = [1 2 3 4 5 6 7 8];
y = [0.39 0.47 0.60 0.21 0.57 0.36 0.64 0.32];
plot(x,y,'ko-') % Correction made here!
xlim([0 9])
xlabel('Session')
ylabel('Score on Test')
```

Output 1.8.2:



The point of this example is that you should avoid being too self-congratulatory, at least until you know you have something to be very proud of. We hope you will reach that point! Be open to negative feedback. In that connection, we authors welcome corrections and suggestions about ways to improve this book. Feel free to contact us. We will welcome constructive comments.

1.9 Programming With a Friend

No matter how open you may be to negative feedback, it is hard to catch all the mistakes you may make. And no matter how useful it may be to talk aloud in forming your plan for a computer program, you may feel uncomfortable speaking to no one in particular, especially when others are in earshot.

A good way to avoid these problems is to have a friend by your side while you program. This is one of the best ways to program, in our opinion. Apart from the fact that the interactions can be fun, having two pairs of eyes and ears on a problem can spur creativity.

We encourage you to program with someone else. The co-authors of this text often share questions and suggest solutions with each other, even though we usually collaborate at a distance. If you are using this book in a course, we encourage your instructor to find ways of grading your work so cooperation with others counts for you, not against you.

1.10 Writing Concise Programs

It is fairly easy to write a program that has many unnecessary variables and superfluous lines. It is harder, at least early in training, to write a program that does the same job with few variables and lines. It becomes a source of pride to programmers when they write concise programs. Such programs do more than appeal to programmers' aesthetic sense. Concise programs also tend to finish in less time than programs that are verbose, go on and on, are redundant, and have far too many words in them, as in this needlessly long sentence that should have ended long ago had we not wanted to make the point that excess verbiage isn't helpful.

Sometimes, but not always, a concise program can reduce the time to run a program by seconds, minutes, hours, or even days. If the program must solve a problem on which people's lives depend, finding a quick solution can literally mean the difference between life and death. In more mundane terms, when a program is used to acquire behavioral data, if it runs too slowly, not all potential data can be captured. That said, it is of course possible to write *too* concisely, so the code is obscure to other readers and maybe even to yourself once you've set it aside for a while. Our advice, then, is to be concise, but only to the extent necessary. Don't obsess about writing code that's ultra-brief if it makes it harder for you or others to understand it.

1.11 Writing Clear Programs

As just said, program conciseness can enhance clarity, but that's not always the case. Just as you should be as lucid as possible about what your program must do (the second principle in the list above), you should write programs that are as easy as possible to read and understand. Program clarity becomes especially important when you have written many programs. If you return to a program that you wrote days or weeks ago and find yourself unable to understand it, you will be very frustrated.

There are several things you can do to make your programs clear. One is to use extra lines of code or extra variables to make the structure of the program transparent. For example, if you need to divide one term by another and the numerator and denominator both contain complex expressions, it usually helps to have one variable for the terms in the numerator and another variable for the terms in the denominator. The quotient can then be expressed as the ratio of the two variables. The program might have a few more variables than are strictly required, but it will be easier for you and others to understand the code later.

A second practice to make your code clear is to give your variables meaningful names. For example, in the program presented earlier (Codes 1.8.1 and 1.8.2), it would have helped to call the independent variable `session` rather than `x` and to call the dependent variable `test_score` rather than `y`. Using those meaningful variable names might have prevented the "accidental" plotting of `x` against `x` rather than the more appropriate plotting of `y` against `x`.

A third practice to improve program clarity is to add comments. Comments are nonexecutable statements that provide information for the programmer (or reader) instead of for the computer. In MATLAB, comments are preceded by a percent sign (%), as shown in Code 1.8.2.

Programmers comment in different ways. Some interleave comments and executable lines of code. Others tend to provide comments above the executable code (at the start of the program), putting relatively few comments in the body of the program. The first author of this book prefers the latter method because it allows him to provide a conceptual plan for the program to follow, along with introductions of the variables he will be using. He prefers not to have too many comments interspersed with code within his programs because he finds them distracting to read and, frankly, a pain to write.

Providing comments at the start of a program can help you start your programming session by combining the need for commenting with the need for “speaking aloud.” Developing a plan for a program is often aided by putting the plan into words, as stated earlier (Section 1.5). Being able to say what your program should do will help you write the code you need. The first author often sits down and starts typing the description of what his program will do, editing the emerging comment until he reaches the point where he thinks the procedure he’s describing is as clear and mechanically doable as he can make it. Then he begins coding, testing one part of the code at a time, saving successive edits in files with higher and higher version numbers.

Here is an example of one such program. The comments in the opening section (before any executable statements) are typical of what the first author writes. In a short program like this, no further comments are usually needed, because once you gain familiarity with MATLAB, the meanings of the executable statements can usually be understood if the context is clear. All the commands used below will be explained in more detail later in this book.

Code 1.11.1:

```
% Largest_So_Far_01

% Find the largest value in the one-row matrix V.
% Initialize largest_so_far to minus infinity.
% Then go through the matrix by first setting i to 1
% and then letting i increase to the value equal
% to the number of elements of V, given by length(V).
% If the i-th value of V is greater than largest_so_far,
% reassign largest_so_far as the i-th value of V.
% After going through the whole array, print out
% largest_so_far.

V = [7 33 39 26 8 18 15 4 0];
largest_so_far = -inf;
for i = 1:length(V)
    if V(i) > largest_so_far
        largest_so_far = V(i);
    end
end
largest_so_far
```

Output 1.11.1:

```
| largest_so_far =  
|      39
```

The foregoing program can be adapted to find the largest value of other arrays, including much larger ones. We include the program here to give you a taste for what MATLAB programs look like. We also want to convey the idea that it's advisable to test programs on small scales. In general, it's advisable to work on "toy" problems before scaling up to larger ones. This program was tested with an array of length 9. Nine numbers is a more tractable length to use at first than 9,000,000. Just to be sure there are no problems, the program should also be tested with sample data sets in which the largest value is in the first or last position of the matrix because many program errors only reveal themselves at such boundaries.

One last point about program clarity follows. Like all writing, a program is composed for several audiences. Apart from yourself (the person writing and using the code), there are three audiences to keep in mind.

First, there is the computer. The computer, the machine, must be able to deal with the program in the way you wish. At the very least, the program supplied to the computer must be syntactically and logically correct.

The second audience is a colleague, who may wish to evaluate or adapt your program for a related purpose. The colleague may need to understand your program and its logic, with or without your direct advice, and without any particular insights into how you addressed the problem beyond the comments you provided.

The third audience is your future self who, another day, may look back at the prior work. At that later time, you may be faced with understanding what you did without a detailed memory of how you addressed the problem. In the urgency of writing your program to solve an immediate problem, you may take shortcuts, such as using very brief mnemonics for variable names, the meaning of which may be forgotten in the future. To ensure against this unhappy outcome, you may find that once the program is completed, it will serve you to spend a little time clarifying the variable names and adding a few judicious comments. Once you have made these changes, be sure to test the program again, lest your clarification inadvertently produced a new error.

In that spirit, here is the program from Code 1.11.1, with the variable name `V` replaced by `theDataArray`. A couple of other variables and comments have been added as well. Is it clearer to read? Is the result different? Try to make your own programs "self-documenting" by selecting variable names and comments that are as self-explanatory as possible.

Code 1.11.2:

```
| % Largest_So_Far_02  
  
| % Find the largest value in the one-row matrix theDataArray.  
| % Initialize largest_so_far to minus infinity.  
| % Then go through the matrix, by first setting i to 1  
| % and then letting i increase to the value equal
```

```
% to the number of elements of theDataArray, given by
% length(theDataArray).
% If the i-th value of theDataArray is greater than
% largest_so_far, reassign largest_so_far with the i-th
% value of theDataArray.
% After having gone through the whole array, print out
% largest_so_far, which will be the largest value found.

theDataArray = [7 33 39 26 8 18 15 4 0];
%start with an absurdly small maximum
largest_so_far = -inf;

for i = 1:length(theDataArray)
    if theDataArray(i) > largest_so_far
        %Got a new candidate!
        largest_so_far = theDataArray(i);
    end
end

% All done...so what's the maximum?
largest_of_them_all = largest_so_far
```

Output 1.11.2:

```
largest_of_them_all =
    39
```

1.12 Writing Correct Programs

If your program does not generate any error messages and generates plausible output, does that mean the results are correct? You will find that the MATLAB programming environment, introduced in Chapter 2, serves as an excellent source of feedback as you write and then try to run your own programs. You will be told, indirectly or directly, if your syntax (word use and punctuation) is acceptable or unacceptable. If your syntax is unacceptable, you will get an error message. Otherwise, your program will run. If you get an error message, it will be up to you to figure out what needs to be done to resolve the error. It takes some time to learn to interpret error messages, but over time you will learn to do so.

If your syntax is acceptable, it is your responsibility to confirm that the output you get is correct, because correct syntax alone does not guarantee correct program logic. You will find that judging the correctness of your program's output is often as challenging as generating acceptable syntax. As in natural language, an expression can be syntactically correct but not mean what you intend. Sometimes a program seems to work, but lurking within it is some subtle error that makes the output obviously wrong or, much worse, seemingly correct but actually flawed.

Detecting such mistakes is one of the most challenging aspects of programming. In general, developing a program that works correctly requires more than an understanding of programming syntax. It also requires greater clarity and explicitness about procedures to

be followed than is usually required in daily life. Additionally, it requires some way of verifying the output. Striving for such clarity and explicitness is one of the things that makes programming a humbling, though educational, experience.

As you plan your program, pay attention to the eventual means of verification as well as the logic of computation. For instance, suppose you have a set of reaction-time data from a within-subjects design experiment. In such a design, the number of trials observed in each condition may be determined by the experimental design. Part of the output of the analysis program that you write can be the number of trials in each condition (`n_trials`) for each subject, even if those values do not enter into subsequent analyses. You can take getting the correct (i.e., predicted) values of `n_trials` in each condition as evidence that all trials have been considered in the analysis. Conversely, any apparent anomaly in the values of `n_trials` may alert you to an error somewhere, whether in data acquisition or in the summary computation.

Relatedly, if a program analyzes the data of dozens of participants, it is well worth performing the analysis of at least one or two participants by hand, if possible, to verify the match between the computer's computations and your own. In fact, beginning by doing one subject by hand will give you insights into how best to approach the programming problem. Similarly, it's not a bad idea to have two researchers each independently write a program to analyze the same data. If the two programmers' results agree in every detail, you can be reasonably confident in the correctness of the analysis. If it turns out that there is some detail in which the two results do not agree, that outcome provides an opportunity to explore the difference to see if it is due to a programming error, a difference in understanding the data, or error(s) in the analysis logic.

Another useful shortcut for data verification is to exploit a different analysis environment to serve as that "second programmer." The results of analyzing a small subset of the data in a spreadsheet or statistical package should agree perfectly with the corresponding output of your MATLAB program. If the results differ, even apparently trivially, you will want to track down the source of the disagreement.

1.13 Understanding How the Chapters of This Book Are Organized

If you are persuaded that it makes sense for you to go further with this book, it will help you to understand how the book's chapters are organized.

Each chapter begins with the sentence, "This chapter covers the following topics," after which those subjects are listed. The way the subjects are listed is via presentation of the chapters' section names. All the section names of this book begin with gerunds, such as "Understanding . . .," "Approaching . . .," or "Deciding . . ." The sections are titled this way because we want you to learn by doing. You should be actively engaged in understanding, approaching, and deciding (to name some activities) as you pursue the material presented here.

Continuing with the layout of the chapters, after all the section titles are given, each chapter continues with the sentence, "The commands that are introduced and the sections in which they are premiered are:." This sentence precedes a list of all the new commands introduced in the chapter, along with the sections in which those new command are first discussed. If you run your finger down the list and find the activity to which it corresponds, you should

be able to turn to that section and find an example of how the command is used. The commands discussed are also listed alphabetically, with reference to their first mention, in the Commands Index.

Every program shown in this book has a code number. The first number (to the left of the decimal point) corresponds to the chapter in which the code appears. The second number (between the two decimal points) corresponds to the section in which the code appears. The third number (to the right of the second decimal point) is the number of the code within the section. All MATLAB code appears in *Courier* font, as do all words taken from the code shown in the text body of this book.

Every program that yields output has its output shown in the same format as the code. The output has a number that corresponds to the code that produced it.

One thing that is missing from the programs shown in this book are extensive comments. We have left them out not because comments are unimportant but because, for most of the programs in this book, the comments are, in effect, presented in the text leading up to the programs. If you imagine percent signs in front of the lines of text preceding the code for a program shown here, you effectively have the kind of comment that can be supplied in a program.

Does it make sense for you to read the code shown in this book? Shouldn't you just dive in code for yourself, sinking or swimming as the case may be? We don't want you to sink. We want you to swim, and we think there is much to be learned by reading successful code to figure out what it does and how it does it. You can learn by example. Starting with examples of code can be one of the best ways to learn to program. You can always edit the working example for your own needs, much as a cook can edit a recipe he or she reads in a cookbook.

1.14 Using the Website Associated With This Book

As you leaf through this book, you will see that all the graphs and images are in grayscale. The programs that yield these graphs and images allow for color graphics. The reason the book has grayscale images is to keep the cost of production down, which translates into a lower price for you. You can see the color images generated by the programs, and animations, by going to the website associated with this book (www.routledge.com/9780415535946). You will be able to copy the programs and outputs as you wish.

1.15 Obtaining and Installing MATLAB

How can you access MATLAB? You or your institution can purchase individual or shared licenses. Students can also purchase the educational version for their own use.

MATLAB is, formally, a cross-platform programming environment with versions for Windows, Mac OS, and Unix. There are superficial differences between the Windows version of MATLAB and the version that runs under the Mac OS or Unix operating system. If a program involves certain kinds of input-output, there may be differences across platforms, but these will not interfere with your mastery of the basics of the language.

The differences between the Windows and Mac OS platforms relate primarily to common platform-specific GUI (graphical user interface) conventions and aspects of interfacing for real-time data acquisition. Most of the computational features of MATLAB are equivalent across platforms, so programs written on one platform should work on another. Where there are important platform differences that can cause problems, we will point them out, though we cannot anticipate all problems that might arise.

As of this writing, we have used versions MATLAB installed in the following contexts:

- As a stand-alone program, individually licensed to a particular researcher under an academic license.
- As a stand-alone program (student version), individually licensed to an undergraduate or graduate student.
- Under an educational site license in which the number of simultaneous users on a campus is monitored by a local server.
- As a program that runs remotely on a central server, to which a limited number of simultaneous users may log on.
- Using an open-source alternative to MATLAB, called OCTAVE (www.gnu.org/software/octave/), that allows the running of much of the code of MATLAB. OCTAVE lacks the closely coordinated debugging and program management tools of MATLAB, and we have found that its graphics are less sophisticated, but it is capable of most of the computational operations of MATLAB.

The examples in this text should almost all run identically regardless of the environment and MATLAB version (“release”) that is used. For the most part, we have relied on the current Windows release, R2013a, released March 1, 2013. Because successive MATLAB releases are upward compatible (later versions are compatible with earlier versions), what you learn here should apply to later releases.

How should MATLAB be installed? It is outside our scope to describe the installation procedures needed to get MATLAB to run wherever you are, in part because the details vary depending on the version you are using, the platform you are running on, and the type of license you hold. Ideally, you will have local knowledge to draw on, but MATLAB support through The MathWorks, Inc., is typically very responsive to calls for installation assistance, provided you have your license number handy; see the `ver` command in Chapter 2, Section 2.2.

1.16 Acknowledging Limits

The final section of this chapter is concerned with the limits of this book, our limits as the book’s authors, and the limits of MATLAB itself. It is important for you to know what these limits are so you won’t form unrealistic expectations.

First, with regard to the book, you should know that you will not be able to program in MATLAB if you just read this book without also trying to program yourself. Reading how

to program is a little like reading how to ride a bike. You have to get on and try it yourself. Don't worry if you fall off a few times. Indeed, experienced as we authors may be, in preparing the examples in this book we had to spend quite a bit of time getting the syntax to work just right, often with many cycles of the edit-run-error-edit loop. It's no reflection on your skills, then, if you have lots of false starts when putting together a new programming project. We, the authors of this book, have gone through those false starts ourselves.

You should also know that the material presented in this book is meant to *acquaint* you with MATLAB but not to convey every aspect of this vast language and its associated applications. This book would be much denser if it went into many more detailed and advanced aspects of the MATLAB programming language. You should be able to delve into these topics on your own having worked through the material provided here.

Third, you should know about the limits of MATLAB. The "word on the street" is that MATLAB is terrific for graphics and for creating conceptual models. Its reputation is less secure when it comes to real-time data gathering, where commercial or free alternatives like E-Prime, PsyScope, and SuperLab are often favored. For large-scale number crunching or statistics, C/C++, R, SPSS, or SAS may be better than MATLAB. On the other hand, MATLAB is being actively enhanced in so many quarters that its limitations, whatever they may be, will probably wane over time as needed tools are being developed to address deficiencies that are spotted by the MATLAB community.

Three examples of such tools can be mentioned here. One is Psychtoolbox (discussed in Chapter 13), which has methods for precise real-time control in psychophysical research. Another tool is an add-on toolbox, MATLAB Coder (not discussed further in this book), which enables MATLAB programs to be converted and distributed as C++ code. A third toolbox from The MathWorks, Parallel Computing, enables intensive computation to be distributed across multiple processors if your computer has more than one. You can learn more about these and other toolboxes provided by The MathWorks by going to their website.

Another comment about the limits of the book is that while the program examples presented here should be comprehensible to you as a behavioral scientist (veteran or fledgling), the program examples are not drawn from a particular approach or finding. The interests of behavioral scientists are highly varied, so the examples offered here are generic rather than specific. They are selected more to highlight particular features of MATLAB than to address specific scientific questions.

2. Interacting With MATLAB

This chapter covers the following topics:

- 2.1 Using MATLAB's windows
- 2.2 Using the Command window
- 2.3 Writing tiny programs in the Command window
- 2.4 Allowing or suppressing outputs by omitting or including end-of-line semi-colons
- 2.5 Correcting errors in the Command window
- 2.6 Writing, saving, and running larger programs as scripts (.m files)
- 2.7 Running and debugging MATLAB programs
- 2.8 Keeping a diary
- 2.9 Practicing interacting with MATLAB

The commands that are introduced and the sections in which they are premiered are:

<code>calendar</code>	(2.2)
<code>clc</code>	(2.2)
<code>ctrl-c</code>	(2.2)
<code>date</code>	(2.2)
<code>disp</code>	(2.2)
<code>doc</code>	(2.2)
<code>exit</code>	(2.2)
<code>help</code>	(2.2)
<code>ls</code>	(2.2)
<code>open</code>	(2.2)
<code>pwd</code>	(2.2)
<code>quit</code>	(2.2)
<code>ver</code>	(2.2)
<code>who</code>	(2.2)
 <code>;</code> (<i>output suppression</i>)	 (2.4)
 <code>up-arrow</code>	 (2.5)
 <code>%</code>	 (2.6)
 <code>...</code>	 (2.6)
<code>commandwindow</code>	(2.6)
<code>ctrl-[</code>	(2.6)
<code>ctrl-]</code>	(2.6)
<code>ctrl-0</code> (<i>zero</i>)	(2.6)
<code>ctrl-i</code>	(2.6)

edit	(2.6)
F5 key	(2.6)
New Script button	(2.6)
Run button	(2.6)
diary	(2.8)
type	(2.8)

2.1 Using MATLAB's Windows

To use MATLAB, you must launch the program. MATLAB is activated, as are most computer applications, by clicking on its icon on the computer desktop or wherever its icon is located. When MATLAB is running, a number of windows will be opened, often as panes docked together in a single window.

When MATLAB is first launched, the **Command** window appears as a pane in the composite window (the one with the name beginning “MATLAB . . .,” followed by the version of MATLAB that you are running). The Command window is the most important window in MATLAB. It is where you control what happens and where you see the results of your programming efforts. The Command window will be described in more detail in Section 2.2.

The second most important window is the **Editor** window, which usually appears as a separate window (the one named “Editor -. . .” followed by the location and name of the file you are editing). Here you exploit MATLAB’s editing capabilities by writing, revising, and saving program scripts and functions, both of which are files that end with a `.m` suffix. The Editor window will be discussed in Section 2.3. Suggestions for how best to arrange these windows will be given in Section 2.5.

The two windows just mentioned are the ones that are most critical. Both are normally used to write and run MATLAB programs. There are also several other windows, however, which are more specialized and are described briefly below.

One is the **Help** window. This window provides a portal to MATLAB’s tutorials. The Help window can be opened directly by entering a command in the search bar at the top right of the MATLAB window, or it can be opened indirectly by typing the `doc` command in the Command window.

The **Command History** window chronicles the commands used in the Command window. You can use this information to remind you what commands you have issued in a MATLAB session.

The **Current Folder** window lists the contents of the working directory. You will learn how to change the Current Folder in Chapter 6 (“Input-Output”). By default, the Current Folder is set to `My Documents/MATLAB` in Windows, and `Documents/MATLAB` in Mac OS.

The **Workspace** window lists the variables that are currently active, giving their names and values. The values of a variable can be viewed in this window in spreadsheet form by clicking on the grid icon to the left of its name.

Other windows, called **Figure** windows, can be created, opened, and closed in your programs to show graphics, text, and other related information (e.g., sounds). Details will be given in Chapter 9 (“Plots”).

2.2 Using the Command Window

As mentioned above, after MATLAB is activated, it brings up the Command window. This is the window where you can issue commands. You do so by typing after the `>>` prompt.

Some useful commands that can be typed after the `>>` prompt are given below, followed by the purposes they serve. It will be helpful for you to read through this list now because the commands are listed more or less “chronologically,” in a way that corresponds to what occurs in a typical MATLAB session. Some of the commands tend to be used more than others. The most frequent ones, in our experience, are `help`, `ls`, `pwd`, `edit`, `open`, `ctrl-c`, and `exit`.

<code>ver</code>	Information about your license, computer, and MATLAB version, together in a convenient summary. If you consult with MathWorks support, you will need this information.
<code>date</code>	The current date (in a format you can specify).
<code>disp</code>	The value of an expression (numeric or string), displayed in the Command window.
<code>calendar</code>	The calendar for the current month.
<code>help</code>	Topics for which help can be provided within the command window. Adding a topic name after <code>help</code> (followed by a space) brings up help about that topic, provided it is known to MATLAB. You can find out what topics are known to MATLAB by first typing <code>help</code> alone. This brings up all the categories for which <code>help</code> is available.
<code>doc</code>	This is a shortcut to the Help window, where all the help that can be viewed in the Command window is available, plus more. The Help navigator can also be accessed via the Help tab at the top of the main MATLAB window.
<code>pwd</code>	Identifies the current directory, the one listed in the Current Folder window, and the default location for saving a script. (<code>pwd</code> stands for “print working directory”.)
<code>ls</code>	Lists the contents of the current directory. Adding just part of a file name after <code>ls</code> (following a space) with an asterisk

	replacing part of the file name causes all the files with that named part to be listed. Thus, <code>ls tim*</code> lists <code>tim_program_01.m</code> , <code>tim_program_02.m</code> , <code>timmy_program_101.m</code> , and <code>timothy.doc</code> , provided these files exist in the current directory. <code>ls tim*.m</code> lists <code>tim_program_02.m</code> , and <code>timmy_program_101.m</code> , but not <code>timothy.doc</code> .
<code>open</code>	Opens a file in the current directory or invokes other programs as needed (e.g., Adobe Acrobat for <code>.pdf</code> files).
<code>who</code>	Lists the names of the currently active variables.
<code>whos</code>	Lists the names of the currently active variables along with their sizes and other attributes.
<code>ctrl-c</code>	Holding down the <code>ctrl</code> key and then pressing the <code>c</code> key interrupts the program that is currently running, provided the Command window is the active window (the window in front of any others that are open). This is very useful when you have “runaway” programs and unwanted data are being spewed on the screen or when you have a program that is running for a long time without any output that you actually want.
<code>clc</code>	Clears the Command window.
<code>exit</code>	Terminates MATLAB.
<code>quit</code>	Runs an optional program called <code>finish.m</code> , whose contents can be customized by the user, then terminates MATLAB, just as <code>exit</code> does.

2.3 Writing Tiny Programs in the Command Window

The preceding list of commands is just a small subset of those that can potentially be typed in the Command window. In fact, the number of possible commands that can be typed in the Command window is infinite, because a series of commands of arbitrary length and complexity can be typed or pasted after the command line prompt (`>>`).

In practice, typing or pasting very long series of commands is not a good idea, however, because the longer and more complex the commands, the greater the chance of error. Once your sequence of commands has grown to a few lines (or is expected to be several lines long), it is better to generate program scripts “off-line” in MATLAB’s Editor. There, the scripts can be saved and modified. We will turn to the Editor in the next section. In this section, setting the stage for what will come when we turn to the Editor per se and to acquaint you with some elementary programming, we will consider a few tiny programs that can be written in the Command window. The rules governing acceptable command syntax are the same whether the commands are typed into the command line “by hand”

or are part of a file in the Editor. Therefore, typing commands into the Command window can be a good way to experiment with getting the syntax right before you add the lines to an edited program.

One of the most fundamental programming tasks is to assign a value to a variable. Suppose you want to assign the number 2 to some variable, arbitrarily called A. This can be done by typing `A = 2` after the command line prompt as follows:

Code 2.3.1:

```
| >> A = 2
```

Output 2.3.1:

```
| A =  
    2
```

The ordering of terms in the assignment is important, as shown below.

Code 2.3.2:

```
| >> 2 = A
```

Output 2.3.2:

```
| ??? 2 = A  
  
Error: The expression to the left of the equals sign is  
not a valid target for an assignment.
```

The error message indicates that, in contrast to mathematics, where an equation means the same thing regardless of whether terms appear to the left or right of the equal sign, order matters in MATLAB. Thus, `2 = A` does not mean the same thing as `A = 2`. Programmers often say “A gets 2” when referring to statements such as `A = 2` to indicate that they are referring to a variable assignment rather than to a conventional mathematical equation.

In MATLAB, variable names, program names, and other file names are case sensitive. Consequently, if you query MATLAB about the value of A, you can get a satisfying, if not terribly exciting, result:

Code 2.3.3:

```
| >> A
```

Output 2.3.3:

```
| A =  
    2
```

If you ask MATLAB about the value of a variable called `a`, which you innocently believe is the same as `A`, you get an error message:

Code 2.3.4:

```
| >> a
```

Output 2.3.4:

```
| ??? Undefined function or variable 'a'.
```

MATLAB can, of course, tolerate lower-case variable names, so it is fine to assign a value to `a`:

Code 2.3.5:

```
| >> a = 3
```

Output 2.3.5:

```
| >> a =  
      3
```

In short, `A` and `a` are different variables. You can even carry out computations using your two variables, assigning a new variable to the result:

Code 2.3.6:

```
| >> My_Difference = a - A
```

Output 2.3.6:

```
| My_Difference =  
      1
```

The last example shows that the name of a variable need not be restricted to a single letter. It can be a string such as `My_Difference`. Spaces are not permitted in variable names or in names of programs or other files. However, spaces are useful for reading meaningful phrases like `My Difference`, so a subscript line can be used as a proxy for the space, as in `My_Difference`. Another method is to use uppercase letters to demarcate the words within compounds. This format is called “camelCase.” (Think of what a camel looks like. The capital letter in the middle of the name is like a hump in the middle of a camel’s back.)

Be aware that variable names cannot start with numbers. Neither can they include special characters (`$`, `%`, `&`, `@`, `-`, `+`, `*`, `/`, `\`, `^`, or the comma). Finally, variable names cannot use special, reserved words for MATLAB, like `if`, `for`, or `end`. You will encounter these reserved words later in this book. You needn’t worry about remembering them at this stage,

nor will you *ever* have to worry about this, for if you assign a variable to a reserved word in MATLAB, you will get an error message, with a helpful pointer (!) underneath the character where MATLAB detected the error:

Code 2.3.7:

```
| >> for = 4
```

Output 2.3.7:

```
| ??? for = 4;  
|  
| Error: The expression to the left of the equals sign is  
| not a valid target for an assignment.
```

This last example illustrates how relying on error messages can help you.

In contrast to the check for misuse of reserved words as variable names, there is no automatic check against the inadvertent use of the name of a built-in MATLAB function or command as a variable name. This can create a very-hard-to-track-down problem, as in the following example.

In computing an average, you might be tempted to write `mean = (5+3+1)/3`. This would assign the value of 3.0 to the variable `mean`. However, if you later tried to use the built-in MATLAB function `mean` on another set of numbers, the operation would either fail or (arguably worse!) succeed but return a plausible, wrong value, in this case 3.0, that you had previously assigned to the variable `mean`. Similarly, if you defined `pi = 22/7`, that value would override `pi`'s built-in value of 3.14159... . So it's good practice to use descriptive variable names to avoid functions and variables that are already defined in MATLAB. If you are in doubt about the name of a variable, as well you should be at this stage of your learning (!), check by using the `help` command. For example, you could type `help mean` or `help pi` at the command prompt (>>). Another strategy is to use variable names that are unlikely to be part of MATLAB's library, such as `David_Mean` or `Jon_Mean` or `Brad_Mean`. When in doubt, there is no harm in checking whether a name you are introducing has special status. You can do so by typing `help` followed by the name you are considering. If MATLAB replies that the name is unknown, you can use it safely.

Must a variable name be supplied for the result of every new computation? The answer is No. When no output variable is declared, MATLAB automatically assigns the output to a variable called `ans`, short for "answer."

Code 2.3.8:

```
| >> a + A
```

Output 2.3.8:

```
| ans =  
| 5
```

You can inquire into the value of `ans` just as you can inquire into the value of any other variable. `ans` has whatever value was most recently assigned to it.

Code 2.3.9:

```
| >> ans
```

Output 2.3.9:

```
| ans =  
      5
```

2.4 Allowing or Suppressing Outputs by Omitting or Including End-of-Line Semi-Colons

Including or not including a semi-colon at the end of a line of code has an important effect. Omitting the semi-colon enables output to the Command window, at least for operations that return a value. Adding a semi-colon suppresses screen output for that operation.

For example, if you type `My_Difference` followed by a semi-colon, the result is initially disappointing.

Code 2.4.1:

```
| >> My_Difference;
```

Output 2.4.1:

```
| >>
```

Apparently nothing happened. In a sense, this is true, because `My_Difference` already had a value assigned to it. That value was the one it took on earlier, via Code 2.3.6.

If you do a new computation and follow it with a semi-colon, you again seem to get the same null effect:

Code 2.4.2:

```
| >> My_Difference_2 = A - My_Difference;
```

Output 2.4.2:

```
| >>
```

However, the computation has, in fact, been carried out, as you can confirm by typing `My_Difference_2` without a semi-colon at the end:

Code 2.4.3:

```
| >> My_Difference_2
```


Output 2.4.3:

```
| My_Difference_2 =  
| 1
```

Suppressing outputs by adding a semi-colon at the end of a command can be useful to prevent the printing in the Command window of a huge list of numbers that goes on “forever.” You can break out of such a salvo of unwanted output by interrupting the program with `ctrl-c`, when the Command window is the active window. However, it is better to get into the habit of adding semi-colons to the ends of lines, removing them when you want to see the results of particular lines. The practice of including and omitting semi-colons at the ends of lines of code in MATLAB is so important that we have devoted an entire section to this one feature of MATLAB.

2.5 Correcting Errors in the Command Window

What if you make an error in the Command window? Recently entered lines can be restored to the Command window using the `up-arrow` key. Hitting the `up-arrow` on the keyboard after the `>>` prompt n times brings you back n command lines. For example, hitting the `up` arrow once restores the most recent command, whereupon it can be executed again after modification or correction if needed.

If a line generates an error, the line can be restored (using one or more `up` arrows), and then edited before being executed again (this time correctly). The same procedure can be used to correct an error, say, three lines back, and then with judicious use of the `up-arrow` key and the mouse, the results can be corrected. A useful shortcut, if the line you would like to restore is several lines back and you know its first letter, is to type its first character followed by the `up-arrow` key one or more times.

Trial-and-error correction is helpful when you’re figuring out how to accomplish a particular operation. Here is an example, copied and pasted from the Command window, with optional comments to show the `up-arrow` key’s effects. What was intended was to add 15 to 13 (not to add 5 to 3) and then to show the results in the Command window:

Code 2.5.1:

```
| >> A = 5;  
| >> B = 3;  
| >> C = A + B; %Oops! no output
```

Output 2.5.1:

```
| >>
```

Code 2.5.2:

```
| >> C = A + B %restored by one up-arrow, then ';' deleted
```

Output 2.5.2:

```
| C =  
      8
```

Code 2.5.3:

```
| >> A = 15;      %restored by 4 up-arrows, then 5 changed to 15  
| >> B = 13;      %restored by 4 up-arrows, then 3 changed to 13  
| >> C = A + B    %Short cut! restored by typing C then one  
|                  up-arrow
```

Output 2.5.3:

```
| C =  
      28
```

2.6 Writing, Saving, and Running Larger Programs as Scripts (.m Files)

As already mentioned, once your program gets complicated, in the sense that it doesn't fit into one or two visible command lines, it makes sense to compose and save the program as a script using MATLAB's Editor and then to run it from the Editor window rather than typing it in line-by-line into the Command window. Composing program scripts off-line lets you work on them incrementally. This means that you can add to them a little at a time after checking that each of the components works as expected. Saving the scripts lets you use them again later, either for the same purpose or for inspiration or reminders for future work.

To open the Editor window to modify an existing file or to create a new file, type `edit` in the command line, or click on the `New Script` button. If the name of a file is put in after `edit` (following a space), that file is opened for editing, provided it's in the current directory. Otherwise, a new file by that name is created in the current directory. By default, the filename is assumed to have the `.m` extension. If no filename is specified, the new, as yet unsaved, file is called `Untitled.m` by default.

When the Editor window is first opened, a blank screen appears. Our recommendation for how to proceed next is to write the name of the program as a comment. A comment is a non-executable string, marked by a percent sign (%) to its left at the start of each line. We usually write the title of a program as a comment, then we immediately select and copy the title alone (i.e., without the % sign), and finally we go to `Save As`, pasting the name of the program into the `Save As` dialog. Here is an example of a small program, saved as `My_Program_01.m` with a few comments. The second comment gives a general account of the purpose of the program. Writing such a comment is advisable.

Code 2.6.1:

```
| % My_Program_01  
| % A program to add two numbers called X and Y.
```

```
X = 10;  
Y = 12;  
  
Z = X + Y
```

Output 2.6.1:

```
Z =  
    22
```

To get the output, we pasted the name of the program into the Save As dialog, hit return, and then were gratified to see that the script was saved as a MATLAB script, so defined by its `.m` suffix. To *run* the program, we clicked on the Run button (the green, right-pointing triangle) in the toolbar atop the Editor window. Another way to run a program from the Editor window is to press the F5 key. Regardless of whether you click on the Run button or press the F5 key, your changes will be automatically saved before the script is run.

You can also run a saved program directly from the Command window by typing or pasting its name (without the `.m` suffix) into the command line. You can do this even when the Editor is not running or when the particular program you want to run isn't active in the Editor window. In other words, a program doesn't have to be active in the Editor window to run, as long as it is stored in the current folder. This will prove important later, when you learn how programs "call" other programs.

When you run a program from the Editor window, you will want to see the results. These appear, for the program illustrated above, in the Command window. You don't want to manually select the Command window to see the results, or at least you'd like to avoid the need to do so all the time. The Command window can be activated automatically by including the command `commandwindow` in your program.

As you gain programming proficiency, you will often find it useful to have both the Editor window and the Command window visible. That way, you can quickly see in the Command window the results of a particular run of your program, and you can be set to make further changes back in the Editor window. You can also open more than one file at the same time in the Editor window and switch between the Editor window and the Command window by clicking on the tab buttons at the top of the Editor window. Having two files open at the same time in the Editor window facilitates comparing them or copying useful code from an old program into a new one.

As mentioned earlier, other windows are used on a more optional basis. These windows, like the Command window and Editor window, are accessible at the top of the MATLAB screen via the Windows tab or Desktop tab. Several of these windows can optionally be combined or "docked" as panes in a larger window, or they can be opened as separate windows ("undocked") by dragging their title bars out of the main MATLAB window. Programmers blessed with multiple (or large) monitors often keep several undocked windows open at a time: the Editor and Command windows, one or more Figure windows, and the Help window, each providing feedback about some aspect of program progress or easy access to programming tools. When they are undocked, you can switch between the Command window and Editor window with the `ctrl-0` and `shift-control-0` keys (that's `ctrl-zero`, not `ctrl-oh`). For most control key combinations in Windows, you can substitute the `command` key for `ctrl` in Mac OS.

A couple of other points are important for using the Editor to write MATLAB scripts. One is that each command is usually limited to a single line of code. Sometimes, however, a command must stretch beyond the visible horizon on your computer screen and you may not want to keep scrolling beyond the right edge to see what's there. To make your code more readable, you can add three dots (. . .) to the end of the line and continue the command on the next line. Note that these are three separate dots, not the typesetter's single-character ellipsis (...).

A second point is that you can use blank lines to group related parts of the program. Here is an example that exploits both blank lines and the three-dot construction.

Code 2.6.2:

```
% Continuation_Illustration

Method_1_Score_1 = 899;
Method_2_Score_1 = 1286;

Method_1_Score_2 = 1018;
Method_2_Score_2 = 1344;

Method_1_Score_3 = 1167;
Method_2_Score_3 = 1389;

Summed_Differences_Between_Method_2_and_Method_1_Scores = ...
    Method_2_Score_1 - Method_1_Score_1 ...
+ Method_2_Score_2 - Method_1_Score_2 ...
+ Method_2_Score_3 - Method_1_Score_3
```

Output 2.6.2:

```
Summed_Differences_Between_Method_2_and_Method_1_Scores =
    935
```

To sum up this last point, MATLAB ignores blank lines, unless they appear directly after a line that continues with three dots.

A third point is that the Editor provides convenient tools for indenting and outdenting lines of code. This lets you see the hierarchical structure of your code, which is useful when you have `for` loops, `if-then` structures, and `while` loops, which will be discussed in Chapter 5. The keystroke combinations that allow for indenting and outdenting are `ctrl-[` and `ctrl-]`. An easy shortcut to automatically indent and outdent the entire program, following its syntax, is to select the entire program, by hitting `ctrl-a`, followed by `ctrl-i`.

2.7 Running and Debugging MATLAB Programs

One of the most challenging aspects of programming is to make sure your program does what you want. It is easy to tell that your program *isn't* doing what you want when your program won't run at all—that is, when the MATLAB compiler gives an error message.

Correcting a program that either will not run or that yields odd results is known as *debugging*. Debugging is one of the most important activities in programming. Ideally, debugging should never be necessary. Each of us fantasizes about being so clear-headed and accurate in our coding that we never make mistakes. But no real person is like this, at least if he or she tries to write programs that are the slightest bit challenging. It's important that you know that everyone who writes programs makes mistakes. Needing to debug is inevitable. The techniques for debugging are so varied that we will devote an entire chapter to this topic later in this book.

What should you know about debugging at this early stage? Different people take different approaches. Some people take advantage of MATLAB's debugging resources (some of which are described in Chapter 14). Others prefer a more homespun approach of developing very small programs or small parts of programs, testing them, and then, after debugging them if necessary, adding new code in small steps, checking to make sure the additions work. Both of these approaches can also be combined.

One piece of advice we can offer is aimed at helping you always move forward, never back. Save successive versions of your programs with unique names. Follow the adage expressed in the American slang expression, "If it ain't broke, don't fix it." Once you have a program that works, save it with a name that distinguishes it from its predecessor. Make sure the predecessor program is still available. Thus, if `Behavior_22.m` works well but you plan to make changes to it, immediately begin the editing process by saving the new program with a new name, such as `Behavior_23.m`. Keep the old version so the "surgery" you are about perform on the code in the new version doesn't "kill the patient." You can always return to `Behavior_22.m` and try again, perhaps in `Behavior_24.m`. Remember, computer storage is cheap. Your time is not.

Here are some other bits of advice related more directly to debugging per se. First, when you get an error message, the message will flag the line number of the first offending command in your code. (Line numbers serve no other function in MATLAB than to count and point to lines. You can't refer to line numbers in your code, in contrast to some other programming languages.) If you click on the error message with the line number in the Command window, the Editor will bring you to the line with the problem, or the line where, due to some other earlier problem, the problem is first noticed. So, for example, if your first executable (non-comment) line is `a = 1` and the second executable line is `c = a + b`, you will get an error message, not because there is anything inherently wrong with the syntax of `c = a + b`, but because an earlier line of code is missing: The value of `b` has not yet been assigned.

Another piece of advice about debugging is that you can use breakpoints. A breakpoint is a "stop sign" that can be put on a line of code to stop the program just before executing that line of code. To insert a breakpoint within an already saved program in the Editor window, click on the dash to the left of a line of code (to the right of the line's number). When the program runs, it will stop at the breakpoint, and you can explore the program state by examining the values of variables and change them if you wish in the Workspace window or the Command window. You can then continue executing the program from that point onward, either one line at a time to monitor its progress or at full speed. Examples of using breakpoints are presented in Chapter 14.

Not everyone uses the breakpoint strategy. For example, the first author of this book asks for the value of a variable by adding its name without a semi-colon afterward, followed by `pause` in the next line. When he runs the program, the variable's value pops up in the command window and the program pauses, at which time he either hits the Enter (return) key to let the program go on, or he hits `ctrl-c` to stop the program and attempt to repair whatever caused things to go awry. If the program works, he often turns the two diagnostic lines (the variable name without a semi-colon and `pause`) into comments (usually by selecting them and hitting `ctrl-r`, or [on Mac OS] `command- $\backslash$`). If at some point he wants to “uncomment” the diagnostic lines, he selects them and hits `ctrl-t`.

2.8 Keeping a Diary

You can keep a record of the text that appeared in the Command window of a MATLAB session by using the `diary` function. When MATLAB is activated, `diary` is off. You can designate the file to which you want a diary to be saved with a command like `diary('My_Program_3_diary.txt')`. You can subsequently turn `diary` off with the `diary off` command. When combined with the `disp` command, the `diary` command is a convenient way to generate a text file of the results of a program, as in the following example.

Code 2.8.1:

```
% My_Program_3
diary('My_Program_3_diary.txt')
a = 1 + 2 + 3 + 4;
b = 1 + 4 + 9 + 16;
disp('The sum and sum of squares of the four integers is:')
disp(a);
disp(b)
diary off
```

Output 2.8.1:

```
The sum and sum of squares of the four integers is:
    10
    30
```

The diary file (in this case, `My_Program_3_diary.txt`) can later be opened using the MATLAB Editor, or programs like Notepad, TextEdit, or Word. Conveniently, it can also be displayed in the Command window using the `type` command, which lists the contents of a text file (`.m` or `.txt`) file in the Command window.

Code 2.8.2:

```
| >> type My_Program_3_diary.txt
```

Output 2.8.2:

```
| The sum and sum of squares of the four integers is:  
    10  
    30
```

2.9 Practicing Interacting with MATLAB

Try your hand at the following exercises, using only the methods introduced so far in this book or information given in the problems themselves. Don't look ahead in the text or look things up on the Internet. Try to solve each problem on your own based on what has been presented here so far.

Problem 2.9.1:

Open MATLAB's Command window and get today's date.

Problem 2.9.2:

In MATLAB's Command window, get this month's calendar.

Problem 2.9.3:

Next, look at the calendar for a year ago this month. Hint: Although nothing you have read in this chapter tells you directly how to do this, there was mention of `help`.

Problem 2.9.4:

Find out what time it is using MATLAB by getting help about `clock`. If you first execute the command `format bank`, the output of `clock` will be most readable.

Problem 2.9.5:

In the Command window, add $2 + 2$, and then observe the `ans`.

Problem 2.9.6:

In the Command window, get the result of adding 4 to `ans`. Looking at the new answer, what does this tell you about MATLAB's "willingness" to redefine values?

Problem 2.9.7:

Use the Editor to write and then save a short program called `My_Program_01` which assigns 1 to `w`. Run the program so the value of `w` displays in the Command window.

Problem 2.9.8:

Save `My_Program_01.m` as `My_Program_02.m` and expand it so after `w` gets 1, `x` gets `w + 1`, and then `y` gets `x - 2`. Add another one-line command that brings up the command window.

Problem 2.9.9:

Debug `My_Program_03` so `b` gets the sum of `a` and 3, `c` gets `b - 2`, `d` gets the product of `b` and `c`, and `e` gets `b` divided by `c`.

```
% My_Program_03
a
b = a + 1
c = b - 2
d = b * c
e = b / c
```

Problem 2.9.10:

Write a program called `My_Program_04` in which Code 2.6.2 is expanded so there is a `Method_1_Score_4` that gets 1267, and a `Method_2_Score_4` that gets 1289, and all scores used in method 1 and 2 are subtracted in the way already established. Use `disp` to generate labeled output.

Problem 2.9.11:

What information do you learn from executing the `ver` command?

3. Matrices

This chapter covers the following topics:

- 3.1 Creating matrices
- 3.2 Specifying elements of matrices
- 3.3 Concatenating matrices
- 3.4 Determining the size of matrices
- 3.5 Transposing or reshaping matrices
- 3.6 Creating matrices with shorthand methods
- 3.7 Checking the status of matrices
- 3.8 Clearing and emptying matrices
- 3.9 Practicing with matrices

The commands that are introduced and the sections in which they are premiered are:

`;` (*matrix row delimiter*) (3.1)

`:` (*series delimiter*) (3.2)

`end` (*variable index*) (3.2)

`length` (3.4)

`size` (3.4)

`'` (*transpose operator*) (3.5)

`reshape` (3.5)

`linspace` (3.6)

`logspace` (3.6)

`ones` (3.6)

`zeros` (3.6)

`whos` (3.7)

`[ ]` (3.8)

`clear` (3.8)

`clear all` (3.8)

3.1 Creating Matrices

Computers store and manipulate matrices of values. We humans typically construe those matrices (or arrays) as representing objects and events of interest to us. In behavioral science, we often let matrices of numbers stand for stimuli, responses, response times, response accuracies, and other relevant items. Because of the importance of matrices for behavioral science, you, as a budding behavioral scientist, will want to know how best to

store and manipulate numerical arrays for your own behavioral science research. Later you will learn how arrays of non-numeric symbols, letters, and other special characters, such as '\$,' '!', and '?', can also be represented in matrices.

A single number, such as the number 1, can be thought of as a very simple matrix—a matrix that has just one entry. Recognizing that single-value arrays are arrays like any other can help you turn that idea around and appreciate that there need not be anything special when it comes to arrays with more than one value. Entire sets of numbers can be represented in matrices with multiple elements, typically in one or more rows and in one or more columns. A matrix with just one dimension—either a single row or a single column—is called a **vector**. While a single number always forms a vector, a vector need not always be a single number. Instead, as just indicated, it can be a set of numbers with several rows but one column, or it can be a set of numbers with several columns but one row.

Conveniently, when using MATLAB, you usually don't have to be overly concerned about distinguishing vectors from matrices. MATLAB can treat the two kinds of arrays equivalently.

To help you get a taste of matrices in MATLAB, assign a multi-element matrix of numbers to a variable called A.

Code 3.1.1:

```
| A = [1, 3, 5, 2, 4, 6]
```

Output 3.1.1:

```
| A =
      1      3      5      2      4      6
```

The matrix A is made up of integers, but a matrix needn't be restricted to integers (whole numbers). Real numbers of any sort can be represented in MATLAB matrices.

Code 3.1.2:

```
| B = [4, .8, -.12, 0, -24]
```

Output 3.1.2:

```
| B =
      4.0000      0.8000     -0.1200         0    -24.0000
```

In both of the preceding examples, commas separated the numbers in the matrix, but spaces may serve the same purpose.

Code 3.1.3:

```
| C = [4 .8 -.12 0 -24]
```

Output 3.1.3:

```
| C =  
    4.0000    0.8000   -0.1200    0 -24.0000
```

C is a 1 by 5 matrix, also written as a 1×5 matrix. The first number (1 in this case) refers to the number of *rows* in the matrix. The second number (5 in this case) refers to the number of *columns*.

A convention used in MATLAB, as in matrix algebra, is that the number of rows in a matrix is reported before the number of columns. For this reason, we often refer to a matrix of size $r \times c$. One easy way to recall the row-then-column order of the subscripts in a matrix is to remind yourself of “Royal Crown,” or RC® Cola. Use some other mnemonic if you prefer.

How can you define a matrix that has more than $r = 1$ row? Here we define a 3×2 matrix—that is, a matrix with 3 rows and 2 columns.

Code 3.1.4:

```
| D = [1 2; 3 4; 5 6]
```

Output 3.1.4:

```
| D =  
     1     2  
     3     4  
     5     6
```

Inspection of the code used to define D shows that a semi-colon (;) indicated row endings. After every two elements, a semi-colon was inserted. This gave us the 3×2 layout we wanted.

As this example shows, the semi-colon has an important function in MATLAB besides suppressing printouts (see Section 2.4). Semi-colons within brackets indicate the ends of matrix rows. You can still use a semi-colon at the end of an assignment to suppress printout, as in this example. The output is not shown because there is no output (no printout of D).

Code 3.1.5:

```
| D = [1 2; 3 4; 5 6];
```

MATLAB is very particular about the layout of a matrix. Every matrix must be rectangular. It must have the same number of columns in every row and the same number rows in every column. What happens if this rule is violated? Let’s tempt fate and see.

Code 3.1.6:

```
| E = [1 2 3; 4 5; 6 7 8];
```

Output 3.1.6:

```

| ??? Error using ==> vertcat
| All rows in the bracketed expression must have the same
| number of columns.

```

The error message appeared because in Code 3.1.6, the variable `E` was assigned a row of 3 columns followed by a row of 2 columns. In this case, MATLAB didn't get past the second row. It balked at the second semi-colon, which came one slot (one column) too soon.

At this point, you might want to slam the book shut and walk away, thinking that you may someday have data sets that don't meet the requirement that all rows have the same number of columns or that all columns have the same number of rows. Fear not, or keep your cool. There are ways around this requirement that we will explain later. If there were none, MATLAB would be used by no one!

3.2 Specifying Elements of Matrices

Having defined a correctly formatted matrix of numbers, such as matrix `D` above, you may want to access values in particular locations within the matrix. Suppose you want to know what the number is in the first row of the first column of `D`. You can find this out as follows:

Code 3.2.1:

```

| D(1,1)

```

Output 3.2.1:

```

| ans =
|      1

```

What are those two numbers in the parentheses after `D`? Each number is an **index**. The first index represents “row 1.” The second index represents “column 1.” In effect, you are asking MATLAB, “What value is in the first row and first column of `D`?”

If you want to know what number occupies row 2, column 1 of `D`, you could write

Code 3.2.2:

```

| D(2,1)

```

Output 3.2.2:

```

| ans =
|      3

```

If you want to know all the values in column 1 for all of the rows of D , you could put a colon (:) in the row position and a 1 in the column position:

Code 3.2.3:

```
| D(:, 1)
```

Output 3.2.3:

```
| ans =  
      1  
      3  
      5
```

Think of the colon as representing “from the beginning to the end” or, in this case, “from the first row to the last row.” Building on this idea, if you want to know all the values in column 2 over all the rows of D , you could put a colon in the row position and a 2 in the column position:

Code 3.2.4:

```
| D(:, 2)
```

Output 3.2.4:

```
| ans =  
      2  
      4  
      6
```

To find all the values in row 1 for all of D ’s columns, you could put a 1 in the row position and a colon in the column position:

Code 3.2.5:

```
| D(1, :)
```

Output 3.2.5:

```
| ans =  
      1      2
```

These examples show that when a colon (:) is inserted at a row or column position, it specifies all the values for that row or column. For this reason, the command $D(:, :)$ is equivalent to the command D .

What if you want to see all the elements of the matrix? A single colon will do the trick. As the following code and output show, the output using a single colon as the index reports

the rows of the first column, then the rows of the second column, all in a single one-dimensional array.

Code 3.2.6:

```
| D(:)
```

Output 3.2.6:

```
| ans =
      1
      3
      5
      2
      4
      6
```

The colon can also be used to represent a subset of the values for a row or column by combining it with values representing the starting and ending values, as in this example.

Code 3.2.7:

```
| E = [1 2 3 4; 5 6 7 8; 9 10 11 12]
| PartOfE = E(2:3,2:4)
```

Output 3.2.7:

```
| E =
      1      2      3      4
      5      6      7      8
      9     10     11     12
| PartOfE =
      6      7      8
     10     11     12
```

Just as the colon is useful for referring to specific elements of a matrix, so too is `end`. To get the element in the last row of the second column, you can use this special value. Using this value frees you from having to know how many rows there are or risking the insertion of the wrong value.

Code 3.2.8:

```
| E(end,2)
```

Output 3.2.8:

```
| ans =  
      10
```

To get the second row of the last column, you can write

Code 3.2.9:

```
| E(2,end)
```

Output 3.2.9:

```
| ans =  
      8
```

To get the second-to-the last value in the second row, you can write

Code 3.2.10:

```
| E(2,end-1)
```

Output 3.2.10:

```
| ans =  
      7
```

Finally, to get all but the last value in the second row of E, you can write

Code 3.2.11:

```
| E(2,1:end-1)
```

Output 3.2.11:

```
| ans =  
      5      6      7
```

3.3 Concatenating Matrices

Matrices can be joined together either by rows or columns. Joining two matrices end-to-end is called concatenation. You can combine the two one-row matrices F and G into the two-row matrix H as follows.

Code 3.3.1:

```
| F = [10 11 12];  
| G = [13 14 15];  
| H = [F; G]
```

Output 3.3.1:

```
| H =
      10      11      12
      13      14      15
```

Notice that the semi-colon in the assignment to H, between the variable names, has the same effect as it did in Code 3.2.7 when placed between numbers inside brackets. The semi-colon indicates that the numbers that follow go into the next row of the matrix. If you omit the semi-colon and replace it with a space or a comma, the result, in this case, is a one-row matrix composed of the concatenation of F and G into one longer row matrix.

Code 3.3.2:

```
| H = [F G]
```

Output 3.3.2:

```
| H =
      10      11      12      13      14      15
```

Concatenating two matrices with different numbers of rows and columns causes problems.

Code 3.3.3:

```
| I = [20 21 22 23 24 25 26];
| J = [H;I]
```

Output 3.3.3:

```
| ??? Error using ==> vertcat
| All rows in the bracketed expression must have the same
| number of columns.
```

On the other hand, even though H and I have different numbers of elements, there is no problem with concatenating them into a one-row matrix:

Code 3.3.4:

```
| K = [H I]
```

Output 3.3.4:

```
| K =
      10      11      12      13      14      15      20      21      22      23      24      25      26
```

If you are dealing with multidimensional matrices, MATLAB offers the `cat` function, which allows you to specify which dimension (rows or columns) to combine. If you recall

the order of rows and columns in referring to a matrix, you will see that `cat` across dimension 1 makes more rows, and `cat` across dimension 2 makes more columns.

Code 3.3.5:

```
| cat_rows = cat(1,F, G)
| cat_columns = cat(2,F,G)
```

Output 3.3.5:

```
| cat_rows =
|      10      11      12
|      13      14      15
| cat_columns =
|      10      11      12      13      14      15
```

3.4 Determining the Size of Matrices

Before concatenating large matrices, it is useful to check the size of each one. The size of a matrix, as mentioned earlier, is its number of rows and columns. So the size of matrix `I` is `[1 7]`; that is, it is a 1×7 matrix. You can find the size of a matrix with the `size` function. (Functions, more generally, will be covered in Chapter 4.)

Code 3.4.1:

```
| size(I)
```

Output 3.4.1:

```
| ans =
|      1      7
```

The size of matrix `K` can be found in the same way, and the output can be assigned to a new variable called, in this instance, `sz_K`. As shown below, the output of `size`, when applied to a two-dimensional matrix of the sort we have been considering (with one or more rows and one or more columns) has two values—the number of rows and the number of columns. Matrices with more than two dimensions can also be created, and the results of the size function applied to them have the corresponding number of values. For more information, use `help size`.

Code 3.4.2:

```
| sizeofK = size(K)
```

Output 3.4.2:

```
| sizeofK =
|      1     13
```

You can assign the number of rows and number of columns identified by the `size` function directly to two elements of a new matrix whose elements can be called, if you wish, `rows` and `columns`:

Code 3.4.3:

```
| [rows columns] = size(K)
```

Output 3.4.3:

```
| rows =  
|      1  
| columns =  
|      13
```

The length of a matrix with just one row is the number of elements in that row. Similarly, the length of a matrix with just one column is the number of elements in that column. More generally, the length of a matrix is the *larger* of its number of rows or columns.

Pay close attention to that last statement, for one of us, your humble first author, was unaware of this fact for a while and got strange outputs as a result. When in doubt about the number of rows and columns in a matrix that you have or may generate computationally, don't rely on `length`. Instead, get the number of rows and columns via `size`. When given the `length` command, MATLAB will happily use the larger of the number of rows or columns in the matrix, which may not be what you want.

Studying the following lines of code can give you a feeling for `size` and `length`. In `JJ`, the number of columns is largest, so `length` reports the number of columns.

Code 3.4.4:

```
| JJ = [1:4;5:8]  
| sizeofJJ = size(JJ)  
| lengthofJJ = length(JJ)
```

Output 3.4.4:

```
| JJ =  
|      1      2      3      4  
|      5      6      7      8  
| sizeofJJ =  
|      2      4  
| lengthofJJ =  
|      4
```

In `JJ`, the number of rows is largest, so `length` reports the number of rows. Again, be careful not to invoke `length` when you're not sure whether a matrix has more rows or columns. It is generally safer to use `size` rather than `length` so you can specify the dimension of interest to avoid confusion.

In the case of `KK`, the number of columns (which is smaller than the number of rows in this case) can be determined by specifying the `size` of the second dimension.

Code 3.4.5:

```
KK = [1 5; 2 6; 3 7; 4 8]
sizeKK = size(KK)
lengthKK = length(KK)
sizeKKcolumns = size(KK,2)
```

Output 3.4.5:

```
KK =
     1     5
     2     6
     3     7
     4     8
sizeKK =
     4     2
lengthKK =
     4
sizeKKcolumns =
     2
```

There is a special case of the size of a matrix, which is that it is empty. It is often useful to start with an empty matrix, by assigning “nothing” to it using the bracket notation (`x = []`). Values can be concatenated to it. An empty matrix is a 0×0 matrix; it has no assigned value, and it is not the same as a variable that has a value of zero.

Code 3.4.6:

```
xempty = []
xemptysize = size(xempty)
xempty = [xempty 1]
xempty = [xempty 2]

xzero = 0
zerosize = size(xzero)
```

Output 3.4.6:

```
xempty =
     []
xemptysize =
     0     0
xempty =
     1
xempty =
     1     2
```

```
xzero =
      0
zerosize =
      1      1
```

3.5 Transposing or Reshaping Matrices

Suppose you have two matrices, J and K , defined as follows.

Code 3.5.1:

```
J = [1 2 3 4]
K = [5; 6; 7; 8]
sizeofJ = size(J)
sizeofK = size(K)
```

Output 3.5.1:

```
J =
      1      2      3      4
K =
      5
      6
      7
      8
sizeofJ =
      1      4
sizeofK =
      4      1
```

Because there are no semi-colons between the values in the assignment of J , the size of that matrix is $[1 \ 4]$. On the other hand, because there *are* semi-colons between the values in K , the size of that matrix is $[4 \ 1]$. If you try to concatenate J and K , you will get an error message.

Code 3.5.2:

```
L = [J; K]
```

Output 3.5.2:

```
??? Error using ==> vertcat
All rows in the bracketed expression must have the
same number of columns.
```

You can get around this problem, if it makes sense to do so, by “turning one matrix around.” More technically, you can *transpose* the matrix so its rows and columns are interchanged. MATLAB lets you transpose a matrix by adding an apostrophe (').

Code 3.5.3:

```
| K'
```

Output 3.5.3:

```
| ans =  
      5      6      7      8
```

Matrices J and K' can now be combined into a two-row matrix:

Code 3.5.4:

```
| L = [J; K']
```

Output 3.5.4:

```
| L =  
      1      2      3      4  
      5      6      7      8
```

If you wish to take the transpose of L, you can do so easily:

Code 3.5.5:

```
| L'
```

Output 3.5.5:

```
| ans =  
      1      5  
      2      6  
      3      7  
      4      8
```

As seen here, the first row becomes the first column, and the second row becomes the second column.

Another way of modifying the arrangement of the elements in a matrix is to reshape the matrix. For example, the 18 cells of a 1×18 matrix can be arranged as either a 3×6 or a 6×3 matrix, using the `reshape` function. The rows of column 1 are filled first. Then the other columns are filled in.

Code 3.5.6:

```
A = [3 6 9 12 15 18 21 24 27 30 33 36 39 42 45 48 51 54;
A1 = reshape(A, 3, 6)
A2 = reshape(A, 6, 3)
```

Output 3.5.6:

```
A1 =
      3      12      21      30      39      48
      6      15      24      33      42      51
      9      18      27      36      45      54
A2 =
      3      21      39
      6      24      42
      9      27      45
     12      30      48
     15      33      51
     18      36      54
```

Finally, the elements of a two- or three-dimensional matrix can be addressed as if the matrix were a one-dimensional array by giving just one value or range for the index:

Code 3.5.7:

```
A2(7)
A2(9)
A2(5:8)
```

Output 3.5.7:

```
ans =
     21
ans =
     27
ans =
     15     18     21     24
```

This means that by specifying the colon for the index, the entire array can be addressed. In the result, columns have been concatenated in left-to-right order.

Code 3.5.8:

```
A2(:)
```

Output 3.5.8:

```
ans =  
     3  
     6  
     9  
    12  
    15  
    18  
    21  
    24  
    27  
    30  
    33  
    36  
    39  
    42  
    45  
    48  
    51  
    54
```

3.6 Creating Matrices With Shorthand Methods

All the matrices shown so far are small. If you need to create a very large matrix, it is tedious to type in all the values by hand. Fortunately, MATLAB provides shorthand methods for creating matrices.

Consider the matrix M.

Code 3.6.1:

```
M = [1 2 3 4 5 6]
```

Output 3.6.1:

```
M =  
     1     2     3     4     5     6
```

An easier way to create the same matrix is as follows:

Code 3.6.2:

```
M = [1:6]
```

Output 3.6.2:

```
M =  
     1     2     3     4     5     6
```

The colon tells MATLAB that you want a range of values, in this case going from 1 to 6.

MATLAB lets you specify the increments for the range of values you want. Suppose you want values from 1 to 6 increasing in steps of .5. This can be achieved, as shown here for a matrix arbitrarily called MM.

Code 3.6.3:

```
| MM = [1:.5:4]
```

Output 3.6.3:

```
| MM =
      1.0000      1.5000      2.0000      2.5000      3.0000
      3.5000      4.0000
```

This example shows that inserting a value followed by a colon between the starting and ending values of a matrix (in this case, .5 :) lets you specify the size of the steps to be taken from the starting value to the ending value.

What was the step size before, when we typed `M = [1:6]`? MATLAB “knew” that the step size was 1. The value of 1 was *implicit*. When no value is given in a matrix definition, MATLAB assumes that the desired step size is 1.

The notion that some values are implicit is a very important one. Often, when using MATLAB, you can find sources of flexibility by considering whether there might be a way of specifying a value that seems to be implicitly assigned. Specific examples will come up later—for example, when we discuss properties of figures and the axes used in graphs (see Chapter 9).

Must all matrices have ascending values? Is there a shorthand way to create matrices that have descending values? Not surprisingly, there is. It entails making the step size, and the step direction, explicit.

Code 3.6.4:

```
| Descending_Matrix = [5:-2:-7]
```

Output 3.6.4:

```
| Descending_Matrix =
      5      3      1     -1     -3     -5     -7
```

As this example shows, a negative step sign, coupled with an ending value that is smaller than the starting value, ensures a matrix with descending values. Be sure that the ending value is the one you want. Otherwise, you can get a surprising or unwanted result.

Code 3.6.5:

```
| s = [5:-6:-3]
```


Output 3.6.5:

```
| s =  
      5      -1
```

The final desired value of -3 does not appear here because you can't get to -3 from 5 in steps of -6 .

Errors like this can arise when you want to create a vector (a matrix with a single row or column) with a desired number of values, as for example, when you want to generate a graph with a desired number of points (see Chapter 8). There is a shorthand way to create such a matrix that will ensure your desired ending value is represented. You can use the `linspace` function.

Code 3.6.6:

```
| s = linspace(5,-3,8);
```

Output 3.6.6:

```
| s =  
      5.0000      3.8571      2.7143      1.5714      0.4286  
 -0.7143 -1.8571 -3.0000
```

The `linspace` command, as used here, indicates that you want `s` to be a vector that runs from 5 to -3 with 8 values in all. As seen above, MATLAB has found a step size that yields the desired vector. The step size is the same throughout the matrix. This explains why `linspace` has the name it does. Elements are linearly spaced when the steps between them are the same.

Another function for generating vectors is `logspace`. As its name implies, `logspace` creates a vector whose elements are spaced logarithmically rather than linearly. To learn what `logspace` does (or to remind yourself later), you can use `help` at the current line of the Command window, just as you can use `help` to learn about other commands:

Code 3.6.7:

```
| help logspace
```

Output 3.6.7:

```
| LOGSPACE Logarithmically spaced vector.  
  LOGSPACE(X1, X2) generates a row vector of 50 logarithmically  
  equally spaced points between decades  $10^{X1}$  and  $10^{X2}$ . If  $X2$   
  is  $\pi$ , then the points are between  $10^{X1}$  and  $\pi$ .
```

```
LOGSPACE(X1, X2, N) generates N points.
For N < 2, LOGSPACE returns 10^X2.

See also Linspace, :.
```

What this is saying is that `logspace` generates N points starting with 10 raised to the $X1$ power up to 10 raised to the $X2$ power. When N is not specified, MATLAB sets N to 50.

To make sure you understand this, generate code to check that MATLAB creates a matrix `sss` that has five values spanning 10^1 to 10^2 . As in any logarithmic series, each element should be a constant multiple of the one before it.

Code 3.6.8:

```
logseries1 = logspace(1,2,5)
```

Output 3.6.8:

```
logseries1 =
    10.0000    17.7828    31.6228    56.2341   100.0000
```

Shorthand methods are also convenient for accessing values within matrices. Suppose you want to see just the even-numbered columns of a matrix, or just the odd ones. Here's the way to do that.

Code 3.6.9:

```
myMatrix = [
    1 3 5 7 9 11 13 15
    2 4 6 8 10 12 14 16]
evenColumns = myMatrix(:,2:2:8)
oddColumns = myMatrix(:,1:2:7)
```

Output 3.6.9:

```
myMatrix =
     1     3     5     7     9    11    13    15
     2     4     6     8    10    12    14    16
evenColumns =
     3     7    11    15
     4     8    12    16
oddColumns =
     1     5     9    13
     2     6    10    14
```

Another kind of matrix you may need is one that is all zeros or all ones. MATLAB provides functions for these purposes.

Code 3.6.10:

```
| myZeros = zeros(3,5)
| myOnes = ones(5,3)
```

Output 3.6.10:

```
| myZeros =
|         0         0         0         0         0
|         0         0         0         0         0
|         0         0         0         0         0
| myOnes =
|         1         1         1
|         1         1         1
|         1         1         1
|         1         1         1
|         1         1         1
```

The `zeros` and `ones` functions are often useful, not just to fill matrices with 0's and 1's but also to predefine memory for the results of subsequent computations.

3.7 Checking the Status of Matrices

Several matrices have been created in the programs listed above. What is their status? It's useful to check which matrices are active. This can be done either by activating the Workspace window (see Chapter 2) or by typing `who` (see Section 2.2) in the Command window. Here is the result of typing `who` after creation of the matrices in Section 3.6 (and no others):

Code 3.7.1:

```
| who
```

Output 3.7.1:

```
| Your variables are:
|
| Descending_Matrix  evenColumns      myZeros
| M                  logseries1      oddColumns
| MM                 myMatrix        s
```

Typing `whos` rather than `who` gives more information about the currently active values. You now get the names of the currently active variables as well as their sizes (i.e., the number of rows and columns in each of their matrices), how much memory they use, and the type of variable they are, as well as any relevant attributes. The amount of memory is given in bytes. A byte is a string of eight bits in computer memory. A bit is a binary digit equal to 1 or 0.

Code 3.7.2:

```
| whos
```

Output 3.7.2:

Name	Size	Bytes	Class	Attributes
Descending_Matrix	1x7	56	double	
M	1x6	48	double	
MM	1x7	56	double	
evenColumns	2x4	64	double	
logseries1	1x5	40	double	
myMatrix	2x8	128	double	
myZeros	3x5	120	double	
oddColumns	2x4	64	double	
s	1x8	64	double	

As seen above, the class of all the variables is `double`. An array of type `double` is a matrix of double-precision numbers, that is, numbers that have 14 significant digits. More information about data types will be given in Chapter 7.

3.8 Clearing and Emptying Matrices

To remove a matrix or other variable, you can `clear` it. Suppose you wish to clear `s`.

Code 3.8.1:

```
| clear s
| whos
```

Output 3.8.1:

Name	Size	Bytes	Class	Attributes
Descending_Matrix	1x7	56	double	
M	1x6	48	double	
MM	1x7	56	double	
evenColumns	2x4	64	double	
logseries1	1x5	40	double	
myMatrix	2x8	128	double	
myZeros	3x5	120	double	
oddColumns	2x4	64	double	

Comparing Output 3.8.1 to Output 3.7.2 shows that `s` is now gone.

You can clear all active variables by writing `clear all`. It's good to get into the habit of writing `clear all` at or near the start of a program to be sure you're working with a clean slate.

To reduce the size of a matrix, you can empty some or all of its cells. The following example shows how you can remind yourself of the contents and size of a matrix—in this case `logseries1`—and then empty its last and next-to-last elements by assigning the null element `[]` to them.

Code 3.8.2:

```
logseries1
size(logseries1)
logseries1(end-1:end) = []
size(logseries1)
```

Output 3.8.2:

```
logseries1 =
    10.0000    17.7828    31.6228    56.2341   100.0000
ans =
     1     5
ans =
    10.0000    17.7828    31.6228
ans =
     1     3
```

You can also empty `logseries1` entirely and check its new size.

Code 3.8.3:

```
logseries1 = []
size(logseries1)
```

Output 3.8.3:

```
logseries1 =
    []
ans =
     0     0
```

Emptying a matrix is not the same as clearing it. Clearing a matrix purges it entirely. After a matrix is emptied by setting it to `[]`, the matrix is active and can be added to in subsequent steps. Indeed, an effective way of defining a new matrix to which values will be added is to set it initially to `[]`, as in the first line of Code 3.8.3, and then to add elements to it, as in this example. Here, each concatenation adds a column to a one-row matrix.

Code 3.8.4:

```
matrix_to_be_added_to = []
matrix_to_be_added_to =[matrix_to_be_added_to 1]
matrix_to_be_added_to =[matrix_to_be_added_to 2]
```

```
| matrix_to_be_added_to =[matrix_to_be_added_to 3]
| matrix_to_be_added_to =[matrix_to_be_added_to 4]
```

Output 3.8.4:

```
matrix_to_be_added_to =
    []
matrix_to_be_added_to =
     1
matrix_to_be_added_to =
     1     2
matrix_to_be_added_to =
     1     2     3
matrix_to_be_added_to =
     1     2     3     4
```

To help convey the spirit of the foregoing code, what just happened is a little like adding one item after another to the back of an initially empty pickup truck. Such a truck, affectionately referred to by its somewhat nerdy owner as `matrix_to_be_added_to`, is shown in Figure 3.8.1 as an aid for future memory. This photograph was taken by one of the authors.



Figure 3.8.1

If a semi-colon were included in each line of Code 3.8.4, before the 1, 2, 3, or 4, each concatenation would add a row to a one-column matrix.

3.9 Practicing With Matrices

Try your hand at the following exercises, using only the methods introduced so far in this book or information given in the problems themselves.

Problem 3.9.1:

Create a matrix called `A` that increases in steps of 1 from 1 up to 1,000.

Problem 3.9.2:

Create a matrix called `B` that decreases in steps of 3 from 333 down to 3.

Problem 3.9.3:

Create a matrix called `C` using bracket notation, and define `C` so the result of `[linspace(1,100,100) - C]` is a row of 100 zeros.

Problem 3.9.4:

Create a matrix called `Even` that has the first 200 positive even integers and another matrix called `Odd` that has the first 200 positive odd integers. Check the size of `Even` and the size of `Odd`, as well as `Even(end)` and `Odd(end)` to make sure the values are correct.

Problem 3.9.5:

Repair the following matrix assignments:

`D` should run from 5 up to 100 in steps of .5

```
D = [5:-.5:100]
```

`E` should run from 5 down to -100 in steps of -.25

```
E = [5,25:100]
```

`F` should have 20 values from 1 to 10 that are logarithmically spaced

```
F, = linspace(-1,10.3,23:This is hard(-:
```

Problem 3.9.6:

Consider matrices `G` and `H`, both of size 3×3 :

```
G = [1 2 3; 4 5 6; 7 8 9]
```

```
H = [11 12 13; 14 15 16; 17 18 19]
```

Replace column 1 of `G` with row 3 of `H` using shorthand notation (see Section 3.6).

Problem 3.9.7:

Consider matrix `I`, defined as

```
I = [1:10;11:20;21:30]
```

Empty the last 5 columns of `I` and call the new matrix `J`. Empty the first 2 rows of `J` and call the new matrix `K`.

Problem 3.9.8:

Create a $I \times 4$ matrix called `L` and a $4 \times I$ matrix called `M`. Then concatenate `L` and `M` to create one matrix called `N` of size $I \times 8$, another matrix called `O` of size $8 \times I$, a third called `P` of size 2×4 , and a fourth called `Q` of size 4×2 .

Problem 3.9.9:

Define 2 matrices, `Jack` and `Jill`, as follows.

```
Jack = [1:3:35]
Jill = [41:3:75]
```

Create a new matrix, `Mary`, by replacing every other cell in `Jack` with the values in the corresponding positions of `Jill`. (Hint: What are the lengths of `Jack` and `Jill`? Start by making a matrix, using shorthand notation, that runs from 2 to that length by 2's).

Problem 3.9.10:

Define a matrix `Up` as follows.

```
start_value = 1
step = 2
last_value = 80
Up = [start_value:step:last_value]
```

Define a new value `Down` that is the mirror image of `Up`. Check the output carefully and make whatever change is needed to ensure exact mirroring of `Up` and `Down`.

Problem 3.9.11:

The matrix `LeftToRight` is a 4×4 matrix. Make an array `RightToLeft` that is the left-right mirror image of `LeftToRight`.

```
LeftToRight = [
16  2  3 13
 5 11 10  8
 9  7  6 12
 4 14 15  1
];
```


4. Calculations

This chapter covers the following topics:

- 4.1 Adding, subtracting, multiplying, dividing, and raising values to a power
- 4.2 Using built-in functions to compute the square root, remainder, absolute value, logarithms, and exponentiation
- 4.3 Ordering calculations
- 4.4 Generating random numbers
- 4.5 Performing statistical calculations to obtain the sum, mean, standard deviation, variance, minimum, maximum, correlation, and least-squares fit
- 4.6 Performing statistical calculations with missing data
- 4.7 Calculating with matrices
- 4.8 Using matrix algebra
- 4.9 Sorting arrays
- 4.10 Rounding values, and finding their floor and ceiling
- 4.11 Generating magic squares and calendars
- 4.12 Practicing calculations

The commands that are introduced and the sections in which they are premiered are:

+	(4.1)
-	(4.1)
*	(4.1)
/	(4.1)
^	(4.1)
abs	(4.2)
exp	(4.2)
i (<i>imaginary number</i>)	(4.2)
log	(4.2)
log2	(4.2)
log10	(4.2)
mod	(4.2)
rem	(4.2)
sqrt	(4.2)
()	(4.3)
rand	(4.4)
randi	(4.4)
randn	(4.4)

randperm	(4.4)
reshape	(4.4)
rng	(4.4)
corrcoef	(4.5)
sum	(4.5)
diff	(4.5)
max	(4.5)
mean	(4.5)
median	(4.5)
min	(4.5)
polyfit	(4.5)
std	(4.5)
var	(4.5)
NaN	(4.6)
nanmax	(4.6)
nanmean	(4.6)
nanmedian	(4.6)
nanmin	(4.6)
nanstd	(4.6)
nansum	(4.6)
nanvar	(4.6)
.*	(4.7)
./	(4.7)
.^	(4.7)
* <i>(for matrices)</i>	(4.8)
/ <i>(for matrices)</i>	(4.8)
^ <i>(for matrices)</i>	(4.8)
cross	(4.8)
dot	(4.8)
sort	(4.9)
sortrows	(4.9)
ceil	(4.10)
fix	(4.10)
floor	(4.10)
round	(4.10)
calendar	(4.11)
magic	(4.11)

4.1 Adding, Subtracting, Multiplying, Dividing, and Raising Values to a Power

In the last chapter you saw how matrices can be created and accessed with MATLAB. In this chapter you will see how MATLAB helps you do calculations.

Addition, subtraction, multiplication, and division work as you would expect:

Code 4.1.1:

```
a = 1;  
b = 2;  
c = a + b    % addition  
d = 1;  
e = c - d    % subtraction  
f = 4;  
g = f * 3;   % multiplication (note the use of the  
             % asterisk, *)  
h = f/g      % division
```

Output 4.1.1:

```
c =  
    3  
e =  
    2  
h =  
    0.3333
```

Raising a value to a power is achieved with the caret character (^):

Code 4.1.2:

```
ii = 2;  
j = 3;  
k = ii^j    % ii raised to the j power
```

Output 4.1.2:

```
k =  
    8
```

Finding the n th root of a value is achieved by raising the value to a fractional power. This is possible because the n th root of a value equals the value raised to the $1/n$ power. Thus, the

square root of a value is equal to that value raised to the $1/2$ power, the cube root of a value is equal to the value raised to the $1/3$ power, and so on.

Code 4.1.3:

```
| m = 64;  
| n = 1/2;  
| p = m^n
```

Output 4.1.3:

```
| p = 8
```

It is possible to raise a value to a power expressed in decimal format. Moreover, the power need not be a rational number (a number equal to the ratio of two integers). The ratio of the circumference to the diameter of a circle, π , is an example of such an irrational number.

Code 4.1.4:

```
| pp = 2^.2415  
| qq = 2^pi
```

Output 4.1.4:

```
| pp =  
|      1.1822  
| qq =  
|      8.8250
```

4.2 Using Built-In Functions to Compute the Square Root, Remainder, Absolute Value, Logarithms, and Exponentiation

MATLAB provides a built-in function for taking the square root.

Code 4.2.1:

```
| q = sqrt(m)
```

Output 4.2.1:

```
| q = 8
```

`rem` returns the remainder after division.

Code 4.2.2:

```
| remPiOver3 = rem(pi, 3)
```

Output 4.2.2:

```
| remPiOver3 =  
0.1416
```

The `rem` function is valuable for determining whether a value is odd or even. If a value is odd, the remainder after division by 2 is 1. If a value is even, the remainder after division by 2 is zero. Here we determine what remains after we divide a variable called `subject_number` by 2. In this case, because `subject_number` happens to be 7, the value of the remainder is 1.

Code 4.2.3:

```
| subject_number = 7;  
| remainder = rem(subject_number, 2)
```

Output 4.2.3:

```
| remainder =  
1
```

As you can imagine, determining whether a subject number is odd or even can be useful in assigning subjects to conditions in behavioral science studies. A typical example is assigning subjects to one group if their numbers are odd or to another group if their numbers are even. Those numbers may, in turn, simply reflect the order in which the individuals happened to sign up for the study.

An operation similar to `rem` is `mod`, which also reports the remainder of the first argument divided by the second. As long as both arguments are positive, `rem` and `mod` return the same (positive) values, and if both arguments are negative, they return the same (negative) values. However, the result of `mod(x, y)` has the sign of `y`, whereas the result of `rem(x, y)` has the sign of `x`.

The `abs` function returns the absolute value of its argument. Taking the absolute value of a number (also known as *rectifying* the number) makes the value positive if it's negative.

Code 4.2.4:

```
| abs([-1 3 -5 7])
```

Output 4.2.4:

```
| ans =  
1      3      5      7
```

`exp` is used to raise the base of the natural logarithms to a desired power. Therefore, `exp` is the inverse of the `log` function, which gives the natural logarithm of a number. The base of the natural logarithms is a special value in mathematics, often denoted e . Here e is raised to the 5th power.

Code 4.2.5:

```
| k = exp(5)
```

Output 4.2.5:

```
| k =  
7.3891
```

What exactly is e ? e equals the limit of $(1+1/n)^n$ as n approaches infinity. The value $e^x = \exp(x)$ has the property that its derivative equals itself. The derivative of a dependent variable with respect to some independent variable is the amount by which the dependent variable changes as a result of an infinitesimal change in the independent variable. Thus, the derivative of position with respect to time is the amount by which position changes with an infinitesimal change of time, otherwise known as instantaneous velocity. The fact that the derivative of e^x is itself e^x makes e a convenient constant for modeling change.

It happens that e can be approximated numerically and so can be calculated with a digital computer to a level of precision that is usually adequate for typical needs in behavioral science.

Code 4.2.6:

```
| exp(1)
```

Output 4.2.6:

```
| ans =  
2.7183
```

If you have to use e frequently in your program, you can assign a value to the variable `e` as a shortcut: `e = exp(1)`. Although e is the traditional symbol for the base of the natural logarithms, it is not a reserved term in MATLAB. So if you choose to assign a value other than 2.7183 to `e`, you can do so.

Code 4.2.7:

```
| e = exp(1)  
| e = 12
```

Output 4.2.7:

```
| e =  
2.7183
```

```
| e =  
| 12
```

As mentioned earlier, the inverse of `exp` is `log`, or the natural logarithm, which is denoted in mathematics as *log*, *ln*, or *log_e*. Having earlier set `k` to `exp(5)` and having gotten the value 7.3891, we can ask what value *e* is raised to in order to get 7.3891. The `log` function serves this purpose.

Code 4.2.8:

```
| log(k)
```

Output 4.2.8:

```
| ans =  
| 5
```

When you give a command like `log(k)`, MATLAB assumes that the base of the logarithm is *e*. When you read technical material and come across a term like *ln x*, the term *ln* usually means “natural logarithm,” or logarithm of *x* to the base *e*. *ln* is not a reserved term in MATLAB.

Two other log functions are available for other bases, namely, `log2` and `log10`. While exponentiation to the base *e* is done using the `exp` command, values can be raised to fractional exponents to the bases 2 and 10 using the `^` (exponentiation) operator. The results in Output 4.2.9 and Output 4.2.10 demonstrate the complementary character of the `log` and `exp` functions, as well as their equivalents in other bases.

Code 4.2.9:

```
| log2of_128 = log2(128)  
| two_tothe_7 = 2^7  
  
| log10of_1000 = log10(1000)  
| ten_tothe_3 = 10^3
```

Output 4.2.9:

```
| log2of_128 =  
| 7  
| two_tothe_7 =  
| 128  
  
| log10of_1000 =  
| 3  
| ten_tothe_3 =  
| 1000
```

Code 4.2.10:

```

| valueof_ln_30 = log(30)
| e_tothe_3point4012 = exp(valueof_ln_30)

| valueof_log2_30 = log2(30)
| two_tothe__4point9069 = 2^ valueof_log2_30

| valueof_log10_30 = log10(30)
| ten_tothe_1point4771 = 10^ valueof_log10_30

```

Output 4.2.10:

```

| valueof_ln_30 =
|     3.4012
| e_tothe_3point4012 =
|     30.0000
| valueof_log2_30 =
|     4.9069
| two_tothe__4point9069 =
|     30.0000
| valueof_log10_30 =
|     1.4771
| ten_tothe_1point4771 =
|     30.0000

```

The use of a base other than e or 2 or 10 is futile if you generalize the syntax of the foregoing examples. Here is an attempt, along with the feedback that follows.

Code 4.2.11:

```

| log5(625)

```

Output 4.2.11:

```

| >> log5(30)
| ??? Undefined function or variable 'log5'.

```

Nevertheless, logarithms to bases other than e , 2, and 10 can be obtained using the following formula:

$$\log_b(x) = \log_a(x) / \log_a(b).$$

For example, how many times must 5 be multiplied by itself to produce 625? $\ln(625)/\ln(5)$ tells you the answer is four times.

Code 4.2.12:

```

| logtobase5of_625 = log(625)/log(5)
| FiveToTheFourth = 5^4

```


Output 4.2.12:

```
| logtobase5of_625 =  
|         4  
| FiveToTheFourth =  
|         625
```

Another important quantity in mathematics, traditionally known as i , is the square root of -1 , denoted by `i` in MATLAB.

Code 4.2.13:

```
| sqrt(-1)
```

Output 4.2.13:

```
| ans =  
|         0 + 1.0000i
```

i is an “imaginary” number because the only way to obtain a negative product such as -1 is to multiply a positive number by a negative number ($-1 = 1 \times -1$). This means that taking the square root of a negative value like -1 cannot be the same thing as taking the square root of a positive value like 1 . Yet i has a geometric interpretation, so even though it is an imaginary number, it is not a number that is silly or nonsensical. And why not? The geometric mean of two variables, a and b , is the square root of their product, so it is meaningful to consider the geometric mean of 1 and -1 . The geometric mean of 1 and -1 is the square root of -1 , or i .

A value that has both a real and an imaginary term is called a complex number. Complex numbers are used widely in mathematics and engineering, and are also used in behavioral science. An advantage of this notation is that complex numbers let you express the location of a point in a plane with a single complex number. For example, the complex number $(1 + 2i)$ defines the location of a point in a plane whose x and y coordinates are 1 and 2 , respectively. The first, real, term of the complex number represents the point’s position along the x -axis. The second, imaginary, term represents its position on the y -axis. Knowing this, it is possible to perform algebraic manipulations using complex numbers. For example, you can easily add and subtract complex numbers.

Code 4.2.14:

```
| imaginary = sqrt(1*-1)  
| complex1 = 2*imaginary  
| complex2 = imaginary + (0 + 3i)  
| complex3 = imaginary - 3  
| complex4 = complex1+complex2-complex3
```

Output 4.2.14:

```

imaginary =
      0 + 1.0000i
complex1 =
      0 + 2.0000i
complex2 =
      0 + 4.0000i
complex3 =
    -3.0000 + 1.0000i
complex4 =
      3.0000 + 5.0000i

```

Unless you set `i` to some other value, MATLAB sets `i` to `sqrt(-1)`. You can set `i` to some other value, but if you do, you must clear the variable (`clear i`) to use the symbol `i` again for calculations involving complex numbers.

Code 4.2.15:

```

clear all
% Review of special numbers other than NaN, namely, pi,
% and i.
% Reminder that the default value of i can be overwritten
% but can then be restored by clearing i
The_Special_Number_Pi = pi
The_Special_Number_Sqrt_Minus_1 = i
i = 10;
i_Redefined = i
clear i
After_Clearing_i = i

```

Output 4.2.15:

```

The_Special_Number_Pi =
      3.1416
The_Special_Number_Sqrt_Minus_1 =
      0 + 1.0000i
i_Redefined =
      10
After_Clearing_i =
      0 + 1.0000i

```

4.3 Ordering Calculations

When you program calculations in MATLAB, you often perform more than one calculation per line of code. It's important to be clear about the ordering of operations. The following example shows that outputs involving the same values and operations depend on how the calculations are ordered.

Code 4.3.1:

```
r = 2;
s = 3;
t = 4;
u = 5;
v = 6;

w(1) = r * s - t ^ u / v;
w(2) = r * s - (t ^ u) / v;
w(3) = r * (s - t ^ u) / v;
w(4) = r * (s - t) ^ u / v;
w(5) = (r * s) - t ^ u / v;
w(6) = (r * s - t) ^ u / v;
w(7) = (r * s - t) ^ (u / v);
w(8) = ((r * s - t) ^ u) / v;
w(9) = r * (s - t ^ u / v);

w' % list w(1) through w(9) in column form
```

Output 4.3.1:

```
ans =

-164.6667
-164.6667
-340.3333
-0.3333
-164.6667
5.3333
1.7818
5.3333
-335.3333
```

As seen above, the outcomes differ depending on whether parentheses are used and how the parentheses are positioned. MATLAB, like other mathematical expressions, has a default hierarchy of calculations. For MATLAB, the ordering is exponentiation first, multiplication and division second, and addition and subtraction third. Even knowing this, it is best to include parentheses to avoid unintended results when many calculations are performed in one line. Parentheses can be embedded within other parentheses, as seen in the definition of $w(8)$ above.

Experienced programmers often type the opening and closing parentheses before typing code between them. This helps avoid “parenthesis orphans,” which have an opening parenthesis without a closing parenthesis or vice versa. Parenthesis orphans yield error messages, as seen below.

Code 4.3.2:

```
| w(9) = r * (s - t ^ u / v;
```

Output 4.3.2:

```
??? w(9) = r * (s - t ^ u/v;
      |
Error: Incomplete or malformed expression or statement.
```

4.4 Generating Random Numbers

In doing simulations and conducting experiments in which you want event sequences to be unpredictable, it is useful to generate random numbers. MATLAB provides several random number generators.

The `rand` function generates uniform random numbers between (and including) 0 and 1. The code that follows shows how to assign uniformly distributed random numbers to the elements of a 4×8 matrix.

Code 4.4.1:

```
| uniform_random_distribution = rand(4,8)
```

Output 4.4.1:

```
| uniform_random_distribution =
      0.6225      0.4709      0.2259      0.3111      0.9049
0.2581      0.6028      0.2967
      0.5870      0.2305      0.1707      0.9234      0.9797
0.4087      0.7112      0.3188
      0.2077      0.8443      0.2277      0.4302      0.4389
0.5949      0.2217      0.4242
      0.3012      0.1948      0.4357      0.1848      0.1111
0.2622      0.1174      0.5079
```

`randi` generates integers that are uniformly distributed between 1 and a specified upper-limit integer defined by the first argument of the call to `randi` (5 in the example below). The second argument is the number of rows of the output matrix (1 below). The third argument is the number of columns of the output matrix (8 below). `randi` generates integers randomly and with replacement, so every integer is equally likely regardless of whether it has already appeared. This allows some values to be repeated, as in the following example, which calls for a 1×8 matrix of integers up to the value of 5.

Code 4.4.2:

```
| uniform_integer_distribution = randi(5,1,8)
```

Output 4.4.2:

```
| uniform_integer_distribution =
      2      2      5      1      4      1      2      3
```

`randn` generates normally distributed random numbers rather than uniformly distributed random numbers, as in the prior two functions. Recall that the frequency distribution (histogram) of normally distributed numbers has a bell shape in which approximately 68% of the values fall within ± 1 standard deviation of the mean. By default, `randn` uses a mean of 0 and a standard deviation of 1. Consequently, the generated numbers tend to be close to the mean of 0, with the exact minimum and maximum being unpredictable. The two arguments of `randn` are the number of rows and columns of the generated matrix.

Code 4.4.3:

```
| normaldistribution = randn(4,8)
```

Output 4.4.3:

```
| normaldistribution =  
|      1.7249   -0.9441   -0.2948    0.1133    0.0619  
| 0.4322   -0.0327   -0.8380  
|   -1.0620    0.0485    1.0637   -1.2334    1.7941  
| 0.1206   -0.1556    0.2336  
|    0.8708   -0.5808    1.1224   -1.0238    0.7657  
| -1.9044    0.8514    0.5481  
|    1.4471    0.3301    1.6000   -0.9096    0.1164  
| 1.1801    0.8001    1.3894
```

The normal distribution has a mean of 0 and standard deviation of 1. You can generate a matrix of normally distributed numbers with a specific mean, `mu`, of 10 and a standard deviation, `stdev`, of 15, by adding 10 to the matrix and multiplying by 15.

Code 4.4.4:

```
| mu = 10;  
| stdev = 15;  
| new_distribution = (normaldistribution * stdev) + mu
```

Output 4.4.4:

```
| new_distribution =  
|    5.1879   10.3649   17.0985    4.1751  -16.7249  
| 8.6445    7.5049  -10.1910  
|   -8.9439   -0.0138   25.9044   18.3431    9.5045  
| -5.8584   -8.7725    1.1299  
|    8.7720    4.9679   34.5157   13.6902    0.2338  
| 12.1307    7.8305    8.9384  
|   43.4989  -12.2061    4.8443   -1.8680    8.8181  
| 29.9599   26.2904   19.0650
```

`randperm` lets you generate a random permutation of a specified number of items. It generates a list from which you can sample without replacement. This is a useful

function for tasks like specifying the order of treatments for participants in a behavioral science experiment. Here we specify the random order of 8 treatments for one subject.

Code 4.4.5:

```
| oneSubjectsOrder = randperm(8)
```

Output 4.4.5:

```
| oneSubjectsOrder =  
    6     2     3     8     7     1     5     4
```

What if you had 32 conditions that you wanted to assign, without replacement, to 8 subjects, each of whom would get a different 4 of the 32 conditions? Here is a way to do this using the `reshape` function, which was introduced in the last chapter.

Code 4.4.6:

```
| r = randperm(32)  
| permutedIntegers = reshape(r,4,8)
```

Output 4.4.6:

```
| permutedIntegers =  
    4     2     9    28     1    25    30    11  
    3    12    14    22     7    29    27    21  
   26    20    13    24    18    23    16    19  
    6    17    32    31    15    10     8     5
```

Once you have the matrix `permutedIntegers`, you can assign the values in the first column to the first subject, the values in the second column to the second subject, and so on.

It is important to keep in mind that “random” numbers generated by MATLAB are not truly random. This deficiency is not unique to MATLAB. It is true of all computer programs. The analysis of random number generators and the quest for “truly random” number generators is a longstanding problem in mathematics and computer science. For practical purposes, however, within MATLAB, random numbers (or quasi-random numbers) are generated from a very long pseudorandom sequence that starts from the same place every time MATLAB is launched. You can reset the random number generator to this same starting place with the command `rng('default')`. The dramatic result of `rng('default')` is shown in Code 4.4.7, which assumes MATLAB has just been launched. First, matrix `a` is generated with `randi(8,1,10)`. Later, matrices `b` and `c` are generated with `randi(8,1,10)`. While `b` is different from `a`, as expected, matrices `a` and `c` are identical. This is very unlikely for two random strings, of course, and only occurs because the random number generator was re-initialized before `c` was generated.

Code 4.4.7:

```
a = randi(8,1,10)
b = randi(8,1,10)
rng('default') % or rand('twister',5489) for earlier
                % releases
c = randi(8,1,10)
```

Output 4.4.7:

```
a =
    1     8     1     7     7     7     1     4
    3     7

b =
    4     8     2     3     2     2     7     5
    5     2

c =
    1     8     1     7     7     7     1     4
    3     7
```

Your response to this might plausibly be to say, “Well, then I’ll simply not use the `rng('default')` command.” That’s fine, except you may fall prey to unforeseeable problems as you program many lines of code after launching MATLAB. A better strategy is to use the command, `rng('shuffle')`. This command uses the current time to determine the starting sequence, so the random sequence will always be different.

In the program below (Code 4.4.8) we build on the foregoing suggestions by “shuffling the deck” in the fourth line, after which we save the current random number state in the variable `currentRandomNumberState`.

Code 4.4.8:

```
rng('default')
d = randi(8,1,10)
rng('default')
rng('shuffle')
currentRandomNumberState = rng;
e = randi(8,1,10)
```

Output 4.4.8:

```
d =
    1     8     1     7     7     7     1     4
    3     7

e =
    6     4     4     6     4     1     8     4
    5     2
```

Finally, you can store where you are in the sequence (the “random number state”) using `currentRandomNumberState = rng`, so at a later time you can start where you left off in the sequence (or run multiple simulations with exactly the same random sequence) by restoring the random number generator with the sequence position that you stored in the variable `s`, by giving the command `rng(currentRandomNumberState)`.

In Code 4.4.8 we saved the current state in `currentRandomNumberState`, so in Code 4.4.9 we can start again at the same place, and replicate the same random sequence we got in `e` in our new matrix `f`, even if MATLAB has been relaunched or the random number generator has been shuffled in the interim. The new sequence `f` is identical to the old sequence `e`, which was generated earlier with the random number generated in the state specified by `currentRandomNumberState`.

Code 4.4.9:

```
rng('shuffle')
rng(currentRandomNumberState);
f = randi(8,1,10)
```

Output 4.4.9:

```
f =
     6     4     4     6     4     1     8     4
     5     2
```

4.5 Performing Statistical Calculations to Obtain the Sum, Mean, Standard Deviation, Variance, Minimum, Maximum, Correlation, and Least-Squares Fit

MATLAB provides several functions for statistics. These deserve special attention because of the importance of statistics in behavioral science.

Here is a short program that illustrates some of MATLAB’s built-in functions that are relevant to statistics. The program computes the sum, mean, median, standard deviation, variance, minimum, and maximum of the matrix `r`. As it happens, `r` is the same matrix we used in Chapter 1 to illustrate the process of finding the maximum value for a matrix.

Code 4.5.1:

```
r = [7 33 39 26 8 18 15 4 0];
sum_r = sum(r)
mean_r = mean(r)
median_r = median(r)
standard_deviation_r = std(r)
variance_r = var(r)
minimum_r = min(r)
maximum_r = max(r)
```


Output 4.5.1:

```
sum_r =  
    150  
mean_r =  
    16.6667  
median_r =  
     15  
standard_deviation_r =  
    13.5277  
variance_r =  
    183  
minimum_r =  
     0  
maximum_r =  
    39
```

When you apply the same functions to a multi-row matrix, MATLAB computes the values on a column-by-column basis. To illustrate, we first generate a 3×5 matrix of integers selected between 1 and 10 using the `randi` command.

Code 4.5.2:

```
r = randi(10,3,5)  
sum_vector = sum(r)  
mean_vector = mean(r)  
median_vector = median(r)  
standard_deviation_vector = std(r)  
variance_vector = var(r)  
minimum_vector = min(r)  
maximum_vector = max(r)
```

Output 4.5.2:

```
r =  
     2     8     1     7     4  
     5    10     9     8     7  
    10     7    10     8     2  
  
sum_vector =  
    17    25    20    23    13  
mean_vector =  
    5.6667    8.3333    6.6667    7.6667    4.3333  
median_vector =  
     5     8     9     8     4  
standard_deviation_vector =  
    4.0415    1.5275    4.9329    0.5774    2.5166
```

```

variance_vector =
    16.3333    2.3333   24.3333    0.3333    6.3333
minimum_vector =
     2     7     1     7     2
maximum_vector =
    10    10    10     8     7

```

Another important statistic in behavioral science is the Pearson product-moment correlation coefficient. MATLAB computes this value with `corrcoef`. For technical reasons, `corrcoef` returns a 2×2 matrix if it is called with two arguments, a 3×3 matrix if called with three arguments, and so forth. To learn about those technical reasons, you can type `help corrcoef` at the MATLAB command line. You normally need only the top right or bottom left value of this 2×2 matrix, as seen below.

Here we specify two vectors, *s* and *t*, that have a perfect negative correlation of -1 . Thus, for each increment in *s* there is a corresponding decrease in *t*. The lengths of *s* and *t* must be the same for the correlation to be computed. In the example, we have taken the top-right value of `correlation_matrix` to see the value of *r*. We could have just as easily used the bottom-left value. The main diagonal of a correlation matrix is always ones.

Code 4.5.3:

```

s = [1:20];
t = [50:-1:31];
correlation_matrix = corrcoef(s,t)
r = correlation_matrix(1,2)

```

Output 4.5.3:

```

correlation_matrix =
     1    -1
    -1     1
r =
    -1

```

4.6 Performing Statistical Calculations With Missing Data

In the last chapter, we urged you not to slam shut this book when you learned that MATLAB requires matrices with equal numbers of rows for all columns and equal numbers of columns for all rows. We feared you might because, if you are a behavioral scientist or a budding behavioral scientist, you probably know that sometimes in behavioral science experiments one ends up with missing data. MATLAB provides a special value, `NaN`, to mark such cases. The value `NaN`, as its name suggests, is “Not a Number.” It is not a literal character, nor is it a string of literal characters (see Chapter 8). Instead, it is a special value in a class by itself, “neither fish nor fowl.”

Any element of a matrix assigned the value NaN is an element not to be included in summary statistics of ordinary data. The mean of an array with any NaN's in it will be NaN, so you must be alert to this possibility. If you have thousands of data points that you worked very hard to collect and there happens to be one empty cell to which NaN has been assigned, you don't want to find out that the mean of all your data is NaN. That same summary value will be returned for any other statistical function you might ask for (e.g., max, sum, or var) if your data contains even one NaN and that portion of the data belongs to the set to which that function is applied. Here is a variant of Code 4.5.2 that illustrates our point.

Code 4.6.1:

```
r = randi(10,3,5);
r(3,3:4) = NaN;
r(1:2,2:3) = NaN;
r
sum_vector = sum(r)
mean_vector = mean(r)
```

Output 4.6.1:

```
r =
     2     NaN     NaN     7     4
     5     NaN     NaN     8     7
    10     7     NaN     NaN     2

sum_vector =
    17     NaN     NaN     NaN    13
mean_vector =
    5.6667         NaN         NaN         NaN    4.3333
```

MATLAB provides a special way of computing statistics when there are NaN values in the mix. We will share that with you in a moment, but first want to mention that sometimes it is useful to compute statistics in the normal way, without that special method, to see whether there are NaN values lurking in your data. Just apply a function such as mean to the data and if it comes back NaN, then there's at least one NaN value inside.

Suppose you know that some NaN values do exist in your data. To apply the mean function or some other statistical function to the data, it is necessary to exclude the missing values from the computation. An expression to compute the mean of X when X has missing values is `mean(X(not(isnan(X))))`. Here is an illustration of the use of this approach.

Code 4.6.2:

```
Data = [1 NaN 4 3 NaN 4]
Data_that_are_not_Nans = Data(not(isnan(Data)))
Mean_of_Data_without_Nans = mean(Data_that_are_not_Nans)
```

Output 4.6.2:

```
Data =
    1    NaN    4    3    NaN    4
Data_that_are_not_Nans =
    1    4    3    4
Mean_of_Data_without_Nans =
    3
```

There is another, easier, way to get statistics from data sets that may have a NaN among non-NaN's. In one of the custom toolboxes that MATLAB offers, the MATLAB Statistics toolbox, there are functions that compute statistics for non-NaN values. These functions are `nanmean`, `nanstd`, `nanvar`, `nansum`, `nanmin`, and `nanmax`. As you might guess, these functions calculate, respectively, the mean, standard deviation, variance, sum, minimum, and maximum of the data to which the functions are applied by excluding any NaN's that happen to be in the data.

The following program illustrates how NaN can be assigned to the elements of a matrix and how statistics can then be obtained from the matrix in a way that omits the NaN values in the computation of summary statistics. For clarity, we use just the `nanmean` and `nanstd` functions on the matrix `r` of Output 4.6.1, though, as indicated above, similar functions exist for `nansum`, `nanmedian`, `nanvar`, `nanmin`, and `nanmax`. Note that if a column contains *only* NaN's, any of these statistical functions will return NaN.

Code 4.6.3:

```
r
Column_Means = nanmean(r)
Column_Standard_Deviations = nanstd(r)
```

Output 4.6.3:

```
r =
    2    NaN    NaN    7    4
    5    NaN    NaN    8    7
   10     7    NaN    NaN    2

Column_Means =
    5.6667    7.0000    NaN    7.5000    4.3333

Column_Standard_Deviations =
    4.0415     0    NaN    0.7071    2.5166
```

4.7 Calculating With Matrices

Earlier in this chapter, you read about addition, subtraction, multiplication, division, and exponentiation for single values. Recall that a single value can be viewed as a 1×1 matrix. MATLAB also lets you carry out calculations with larger matrices, as illustrated here.

Code 4.7.1:

```
| u = [1:6]
| v = u + 20
```

Output 4.7.1:

```
| u =
|      1      2      3      4      5      6
| v =
|      21     22     23     24     25     26
```

Here, 20 was added to each element of `u`. A number can also be subtracted from a matrix.

Code 4.7.2:

```
| w = v - 20
```

Output 4.7.2:

```
| w =
|      1      2      3      4      5      6
```

A matrix can be multiplied by a number.

Code 4.7.3:

```
| x = w * 2
```

Output 4.7.3:

```
| x =
|      2      4      6      8     10     12
```

A matrix can be divided by a number.

Code 4.7.4:

```
| y = x / 2
```

Output 4.7.4:

```
| y =
|      1      2      3      4      5      6
```

A number can be added to each element of a multi-row matrix.

Code 4.7.5:

```
| z1 = [1:6;7:12]
| z2 = z1 + 2
```

Output 4.7.5:

```
| z1 =
|      1      2      3      4      5      6
|      7      8      9     10     11     12
| z2 =
|      3      4      5      6      7      8
|      9     10     11     12     13     14
```

When two matrices are added, the elements in corresponding positions are summed.

Code 4.7.6:

```
| z3 = z1 + z2
```

Output 4.7.6:

```
| z3 =
|      4      6      8     10     12     14
|     16     18     20     22     24     26
```

The same holds for subtraction.

Code 4.7.7:

```
| z4 = z1 - 2
| z5 = z1 - z2
```

Output 4.7.7:

```
| z4 =
|     -1      0      1      2      3      4
|      5      6      7      8      9     10
| z5 =
|     -2     -2     -2     -2     -2     -2
|     -2     -2     -2     -2     -2     -2
```

Multiplication, division, and exponentiation (the `*`, `/`, and `^` operators) work on entire matrices, following the rules of matrix algebra (see below). If, instead, you want to apply such an operator on an *element-by-element* basis, as was just done with the `+` and `-` operators, the operator is preceded by a dot. The `.*` operator multiplies matrices element by

element, allowing you to take the products of the values in corresponding row and column positions.

Code 4.7.8:

```
aa = [1:4;5:8]
bb = [4:-1:1;8:-1:5]
cc = aa .* bb
```

Output 4.7.8:

```
aa =
     1     2     3     4
     5     6     7     8
bb =
     4     3     2     1
     8     7     6     5
cc =
     4     6     6     4
    40    42    42    40
```

Likewise, the `./` operator divides element-by-element.

Code 4.7.9:

```
dd = aa ./ bb
```

Output 4.7.9:

```
dd =
    0.2500    0.6667    1.5000    4.0000
    0.6250    0.8571    1.1667    1.6000
```

Similarly, the `.^` operator raises each element of a matrix to an exponent.

Code 4.7.10:

```
dd = aa .^ .25
```

Output 4.7.10:

```
dd =
    1.0000    1.1892    1.3161    1.4142
    1.4953    1.5651    1.6266    1.6818
```

However, for multiplication and division, as noted above, you can scale a matrix (multiply or divide by a single value) using `*` or `/`, without dot notation.

Code 4.7.11:

```
dd3 = dd * 3
halfdd = dd/2
```

Output 4.7.11:

```
dd3 =
    3.0000    3.5676    3.9483    4.2426
    4.4859    4.6953    4.8798    5.0454
halfdd =
    0.5000    0.5946    0.6581    0.7071
    0.7477    0.7825    0.8133    0.8409
```

Scaling a matrix of ones using MATLAB's `ones` function makes it easy to initialize a matrix to some constant value or to all NaN's.

Code 4.7.12:

```
allFives = 5 * ones(2,7)
allNans = NaN * ones(2,5)
```

Output 4.7.12:

```
allFives =
     5     5     5     5     5     5     5
     5     5     5     5     5     5     5
allNans =
    NaN    NaN    NaN    NaN    NaN
    NaN    NaN    NaN    NaN    NaN
```

A particularly useful operation is `diff`, which computes the approximate derivative of a vector by returning the difference between successive items (second minus the first, third minus the second, etc.) in a vector that is one item shorter than the original. To illustrate, we apply `diff` to $y = x^2$ for $1 < x < 8$ to get `d1y`. Then we apply `diff` to `d1y` to get `d2y`. Finally, we apply `diff` to `d2y` to get `d3y`.

Code 4.7.13:

```
x = [1:8];
y = x.^2
d1y = diff(y)
d2y = diff(d1y)
d3y = diff(d2y)
```


Output 4.7.13:

y =	1	4	9	16	25	36	49	64
d1y =	3	5	7	9	11	13	15	
d2y =	2	2	2	2	2	2		
d3y =	0	0	0	0	0			

Students of calculus will recognize $d1y$ as analogous to the first derivative of x^2 , $d2y$ as the second derivative of x^2 , and $d3y$ as the third derivative of x^2 . When position is differentiated with respect to time, the first derivative is velocity, the second derivative is acceleration, and the third derivative is jerk. (A joke that will makes sense to those who are familiar with the breakfast cereal Rice Krispies is that the fourth, fifth, and sixth derivatives are snap, crackle, and pop.)

4.8 Using Matrix Algebra

MATLAB lets you perform calculations that take advantage of matrix algebra. In fact, the word MATLAB is shorthand for “Matrix Laboratory.”

Matrix algebra may be unfamiliar to those behavioral scientists whose education or interests may not have brought them to this subject. If you are in that camp, you can take comfort in the fact that MATLAB provides a medium for exploring more advanced matrix-algebraic operations than the ones we have already covered.

It is not feasible for us to teach matrix algebra here. On the other hand, if you are familiar with it and have grasped the material already presented in this text, you should have little trouble learning the many ways that MATLAB can be used to perform the full range of calculations that are possible in matrix algebra by typing `help *`.

We will illustrate one application of matrix algebra here, just to show its power, which will already be known to those with prior training in these matters, but may prove interesting for those who don’t but are mathematically adventurous. If you are not particularly mathematically adventurous at the moment, you can safely skip to the next section. The remainder of this section is quite technical. In fact, it is the most technical material in this book.

Suppose you want to rotate a vector in a plane. Matrix algebra operations used to do this are similar to other operations in MATLAB. Rotating a vector in a plane is an operation that might be called for in fields like motor control, where it can be useful to compute the postures of a participant reaching for a target. A number of statistical computations, such as factor analysis, also use vector rotations, because the cosine of such rotations is a convenient way to represent the correlations between variables. For our example, we will make the rotations explicit graphically.

Consider a unit circle (with origin $[0,0]$) with a radius that we wish to rotate counter-clockwise by 30° (or $\pi/6$ radians). We represent the X and Y coordinates of the end of that initial radius, $[1,0]$, as a vector, `originalradiuspoint = [1, 0]`. This means

the end of the original radius is one unit to the right of the origin, and zero units from the baseline, so it is a horizontal unit radius pointing to the right (the conventional origin for the radius of a circle in polar coordinates), as shown in Figure 4.8.1. The rotation matrix for rotating any vector counterclockwise around the origin by some angle, θ (expressed in radians) is

$$R = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$$

So, to rotate a radius by 30° , or $\theta = \pi/6$ radians, the rotation matrix is $R = [\cos(\pi/6) \ -\sin(\pi/6); \sin(\pi/6) \ \cos(\pi/6)]$. By convention, positive rotations are counterclockwise around a circle, so values of θ less than zero rotate the point clockwise instead of counterclockwise. We then multiply `originalradiuspoint` by this rotation matrix and see that the endpoint of the new radius, `radiusrotatedOnce`, is at $[0.866, 0.5]$. Similarly, multiplying `radiusrotatedOnce` four more times by the same rotation matrix rotates the radius by 30° four more times (a total of $5\pi/6$ rad, or 150°) so the point of `radiusrotatedFourMoreTimes` now lies on $[-0.866, 0.5]$. The second rotation is accomplished, in the last line of Code 4.8.1, by multiplying `radiusrotatedOnce` by the rotation matrix for the 30° rotation four times, using the exponentiation operator ($R^4 * \text{radiusrotatedOnce}$). The original rotation of 30° , plus the four additional rotations of 30° result in a total rotation of 150° . In Section 9.14 we will see how to represent these vectors graphically. Note that the order of matrix multiplications matters. $[R*a]$ is not the same thing as $[a*R]$.

Code 4.8.1:

```
R = [
    cos(pi/6) -sin(pi/6)
    sin(pi/6)  cos(pi/6)
]
originalradiuspoint = [1;0]
radiusrotatedOnce = R * originalradiuspoint
radiusrotatedFourMoreTimes = R^4 * radiusrotatedOnce
```

Output 4.8.1:

```
R =
    0.8660    -0.5000
    0.5000     0.8660
originalradiuspoint =
     1
     0
radiusrotatedOnce =
    0.8660
    0.5000
radiusrotatedFourMoreTimes =
   -0.8660
    0.5000
```




Figure 4.8.1

Finally, we can express the orientation of any vector such as `radiusrotatedFourMoreTimes` in complex notation (Section 4.2) as follows.

Code 4.8.2:

```
radiusrotatedFourMoreTimes_Complex = ...
radiusrotatedFourMoreTimes(1) + ...
radiusrotatedFourMoreTimes(2) * i
```

Output 4.8.2:

```
radiusrotatedFourMoreTimes_Complex =
-0.8660 + 0.5000i
```

Having just rotated a vector through an angle, we now turn to the complementary function in matrix algebra, computing the angle of rotation between two vectors that have a common origin. We use two of the values that resulted from the prior example, the two radii (`originalradiuspoint` and `radiusrotatedOnce`) which have endpoints at $[1, 0]$ and $[0.8660, 0.5]$, respectively. The `dot` function computes the “dot product” (sometimes called the “scalar product”) of two vectors. The result of this function represents the cosine of the angle between the two vectors (i.e., how much rotation occurred from one to the other). Taking the inverse cosine (using the `acos` function) of the dot product returns

the amount of rotation (in radians), which can be converted to the amount of rotation in degrees by multiplying by $180/\pi$.

Code 4.8.3:

```
vectorR1 = [1,0]; % originalradiuspoint
vectorR2 = [0.8660, 0.5]; % radiusrotatedOnce
dotR1R2 = dot(vectorR1, vectorR2)
RotationR1R2_radians = acos(dotR1R2)
RotationR1R2_degrees = RotationR1R2_radians *180/pi
```

Output 4.8.3:

```
dotR1R2 =
    0.8660
RotationR1R2_radians =
    0.5236
RotationR1R2_degrees =
    30.0029
```

Similarly, for the final vector in the example of Code 4.8.1, we can compute the rotation of the radius from `originalradiuspoint` to `radiusrotatedFourMoreTimes`.

Code 4.8.4:

```
vectorR1 = [1,0]; % originalradiuspoint
vectorR3 = [-0.8660, 0.5]; % radiusrotatedFourMoreTimes
dotR1R3 = dot(vectorR1, vectorR3)
RotationR1R3_radians = acos(dotR1R3)
RotationR1R3_degrees = RotationR1R3_radians *180/pi
```

Output 4.8.4:

```
dotR1R3 =
   -0.8660
RotationR1R3_radians =
    2.6179
RotationR1R3_degrees =
   149.9971
```

The small deviation of Output 4.8.3 and Output 4.8.4 from the “ideal” result of exactly 30° and 150° is attributable to rounding error. We typed in only four significant figures for the endpoints of `vectorR2` and `vectorR3`.

Finally, a related operation in matrix algebra is to compute the axis about which the rotation of a vector takes place. Thinking for a moment in three-dimensional space, and considering a tabletop to be the x-y plane, `vectorR1` and `vectorR3` of the last example have the x-y-z coordinates $[1,0,0]$ and $[-0.8660,0.5,0]$. The zero values for the z-axis in each case simply denote that the vectors are exactly in the x-y plane, (i.e., flat on the tabletop). The

`cross` function computes the “cross product” (sometimes called the “vector product”) of two vectors. The result of this function is a vector whose orientation denotes the axis about which one vector has to rotate to get to the position of the other. The vector’s magnitude is equal to the area of the parallelogram that the vectors span.

Code 4.8.5:

```
vectorR1 = [1,0,0];  
vectorR3 = [-0.8660, 0.5,0];  
AxisOfRotation = cross(vectorR1,vectorR3)
```

Output 4.8.5:

```
AxisOfRotation =  
0 0 0.5000
```

The value of `AxisOfRotation` represents a vector pointing upward from the tabletop (the x-y plane), that is, perpendicular to the tabletop. Thus, the cross product demonstrates that when a vector rotates in the x-y plane, the axis of rotation is along the z-axis (think of the two hands of a clock, and their axis of rotation). The magnitude, .5, of the vector along the z-axis shows that if the parallelogram of which the two vectors form adjacent sides were to be completed by drawing the other two sides, it would have an area of 0.5 units.

4.9 Sorting Arrays

It is often useful to sort values and you can do so with the `sort` function.

Code 4.9.1:

```
r = [3 1 2]  
sorted_r = sort(r)
```

Output 4.9.1:

```
r =  
3 1 2  
sorted_r =  
1 2 3
```

For a matrix with more than one column, you can sort several columns with a single command. Here we sort a matrix based on two sets of random numbers.

Code 4.9.2:

```
rr = [randperm(10)' randperm(10)']  
srr1 = sort(rr)
```

Output 4.9.2:

```

rr =
    10     9
     9     4
     5     5
     1     2
     4    10
     2     6
     7     7
     8     8
     6     1
     3     3

srr1 =
     1     1
     2     2
     3     3
     4     4
     5     5
     6     6
     7     7
     8     8
     9     9
    10    10

```

Note that both columns are now in ascending order, so the original correspondence between the items in each row of matrix `rr` has been lost. You can also sort by one column at a time, however, to retain that correspondence, using `sortrows`. The first argument specifies the matrix to sort. The second argument indicates which column is key. The sign of the second argument denotes whether to sort in ascending or descending order. Here, both columns of `rr` are sorted in ascending order using column 1 as the key, yielding `srr2`. Then both columns are sorted in descending order of column 2, yielding `srr3`. Note the row entries are still paired as they were originally, in `rr`.

Code 4.9.3:

```

srr2 = sortrows(rr,1)
srr3 = sortrows(srr2,-2)

```

Output 4.9.3:

```

srr2 =
     1     2
     2     6
     3     3
     4    10
     5     5
     6     1

```

```
      7      7
      8      8
      9      4
     10      9
srr3 =
      4     10
     10      9
      8      8
      7      7
      2      6
      5      5
      9      4
      3      3
      1      2
      6      1
```

Suppose you need to sort by columns rather than rows. This is easily done use the `sortrows` function combined with the transpose operator, `'`, applied twice, once before sorting and once after sorting to restore the matrix to its original orientation. We illustrate this procedure in the following program, where we have six subjects, each of whom has eight data values, in this case simulated with the command `randi(8,6)`. The matrix `mydata` has the six subject numbers in its first row in the random order than `randperm(6)` provided. We would like the data to be presented with each column having each subject's data but with the order of columns going from subject 1 as the first column up to subject 6 in the sixth column. We do this by transposing `mydata` and then sorting it by its first row in ascending order.

Code 4.9.4:

```
subject = randperm(6)
thedata = randi(8,6);
mydata = [subject; thedata]
mydataSortedbySubject = [sortrows(mydata',1)]'
```

Output 4.9.4:

```
subject =
      2      3      4      1      6      5
mydata =
      2      3      4      1      6      5
      6      3      5      4      7      7
      5      7      4      1      1      5
      6      3      8      1      6      1
      4      4      7      5      8      6
      3      2      6      7      6      5
      4      2      7      7      5      7
mydataSortedbySubject =
      1      2      3      4      5      6
      4      6      3      5      7      7
```

1	5	7	4	5	1
1	6	3	8	1	6
5	4	4	7	6	8
7	3	2	6	5	6
7	4	2	7	7	5

4.10 Rounding Values, and Finding their Floor and Ceiling

MATLAB lets you round down to the nearest integer if the value to the right of the decimal point is less than or equal to .5, and up to the nearest integer if the value to the right of the decimal point is greater than .5.

Code 4.10.1:

```
dd = [
    1.0000    1.1892    1.3161    1.4142
    1.4953    1.5651    1.6266    1.6818]
round(dd)
```

Output 4.10.1:

```
dd =
    1.0000    1.1892    1.3161    1.4142
    1.4953    1.5651    1.6266    1.6818
ans =
     1         1         1         1
     1         2         2         2
```

MATLAB also lets you truncate to the next lowest integer regardless of the value to the right of the decimal point. The relevant function is `floor`.

Code 4.10.2:

```
| floor(dd)
```

Output 4.10.2:

```
ans =
     1         1         1         1
     1         1         1         1
```

You can raise values to the next highest integer regardless of what number appears to the right of the decimal point by using the `ceil` function.

Code 4.10.3:

```
| ceil(dd)
```


Output 4.10.3:

```
| ans =  
|      1      2      2      2  
|      2      2      2      2
```

You can bring values to the next closest integer toward zero regardless of what number appears to the right of the decimal point by using the `fix` function. Here, `fix` is applied both to `dd` and `-dd`. Meanwhile, `floor` is applied to `-dd` to show how the output differs when `floor` or `fix` is applied to negative values.

Code 4.10.4:

```
| fix_dd = fix(dd)  
| fix_minus_dd = fix(-dd)  
| floor_minus_dd = floor(-dd)
```

Output 4.10.4:

```
| fix_dd =  
|      1      1      1      1  
|      1      1      1      1  
| fix_minus_dd =  
|     -1     -1     -1     -1  
|     -1     -1     -1     -1  
| floor_minus_dd =  
|     -1     -2     -2     -2  
|     -2     -2     -2     -2
```

To summarize the effects of `floor`, `fix`, `round`, and `ceil`, here is code used to show a table of their effects on negative and positive numbers.

Code 4.10.5:

```
| disp('          a    floor(a)    fix(a)    round(a)    ceil(a)');  
| a = (-2:.25:2)';  
| b = [a floor(a) fix(a) round(a) ceil(a)];  
| disp(b)
```

Output 4.10.5:

```
|      a    floor(a)    fix(a)    round(a)    ceil(a)  
| -2.0000    -2.0000    -2.0000    -2.0000    -2.0000  
| -1.7500    -2.0000    -1.0000    -2.0000    -1.0000  
| -1.5000    -2.0000    -1.0000    -2.0000    -1.0000  
| -1.2500    -2.0000    -1.0000    -1.0000    -1.0000
```

-1.0000	-1.0000	-1.0000	-1.0000	-1.0000
-0.7500	-1.0000	0	-1.0000	0
-0.5000	-1.0000	0	-1.0000	0
-0.2500	-1.0000	0	0	0
0	0	0	0	0
0.2500	0	0	0	1.0000
0.5000	0	0	1.0000	1.0000
0.7500	0	0	1.0000	1.0000
1.0000	1.0000	1.0000	1.0000	1.0000
1.2500	1.0000	1.0000	1.0000	2.0000
1.5000	1.0000	1.0000	2.0000	2.0000
1.7500	1.0000	1.0000	2.0000	2.0000
2.0000	2.0000	2.0000	2.0000	2.0000

4.11 Generating Magic Squares and Calendars

Many tutorials about MATLAB feature the “magic square”—an $n \times n$ matrix of consecutive integers with the fortuitous, if not truly magical, property that the elements in every row, every column, and both diagonals sum to the same value. MATLAB provides a function called `magic` for generating such matrices.

Code 4.11.1:

```
n = 4;
magic(n)
```

Output 4.11.1:

```
ans =
    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1
```

How useful `magic` will be in your everyday work is an open question. A built-in function that may be more useful is `calendar`. The `calendar` command by itself gives the calendar for the present month. To get the calendar for a specific year and month, say, July 1776, the syntax is as follows:

Code 4.11.2:

```
calendar(1776,7)
```

Output 4.11.2:

Jul 1776						
S	M	Tu	W	Th	F	S
0	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31	0	0	0
0	0	0	0	0	0	0

4.12 Practicing Calculations

Try your hand at the following exercises, using only the methods introduced so far in this book or information given in the problems themselves.

Problem 4.12.1:

You have done a study on the effects of five different treatments on the performance of two groups of participants. One group had earlier exposure to the task, causing their mean score to be 15 points higher than for the first-time group. The data for the two groups are as follows:

```
First_Time_Group = [71 78 80 86 91]
Second_Time_Group = [86 91 97 97 110]
```

What single line of code will remove the 15-point advantage for the `Second_Time_Group`?

Problem 4.12.2:

Continuing with the study of the two groups, and using the original values of `First_Time_Group` and the transformed value of `Second_Time_Group`, compute the mean and standard deviation of `First_Time_Group`, the mean and standard deviation of `Second_Time_Group`, and the mean and standard deviation of the paired differences between `First_Time_Group` and `Second_Time_Group`. Use variable names that will make it easy to understand the output.

Problem 4.12.3:

Assign random permutations of eight treatments, numbered 1 to 8, to each of four participants.

Problem 4.12.4:

Amy participated in a gymnastics competition. She received the following scores in each of four events.

```
Vault: 8.9, 8.7, 8.2, 9.1, 9.0
Uneven_bar: 9.5, 9.3, 9.3, 9.25, 8.9
```

Balance_beam: 8.9, 8.9, 8.7, 8.6, 8.5
 Floor: 8.9, 8.8, 8.8, 8.7, 8.9

Amy's final score for any given event needs to be the mean score after removing the lowest and highest score for that event. Write a program that computes Amy's final score in each apparatus and then the total of all her final scores.

Problem 4.12.5:

You are preparing stimuli for an experiment and discover a mistake in the final column of values on which the stimuli will be based. Each value in the final column needs to be squared. Write a program to correct the error, using just one command. Don't just square each mistaken value by hand. The data that need to be corrected are as follows:

```
Data_Needing_Correction = [23 24 5; 34 35 6; 46 47 7]
```

Problem 4.12.6:

Use MATLAB to verify that the right column of `magic(3)` sums to 15, and that each diagonal of `magic(3)` sums to 15.

Problem 4.12.7:

Earvin "Magic" Johnson wore jersey number 32 when he played basketball for the Los Angeles Lakers. Verify that the top row and rightmost column of `magic(32)` both sum to the same number. Can you think of an easy way to check the sum of all the columns with one command? Likewise for the sum of all the rows? How about checking the main diagonal? (Hint: type `doc diag` in the Command window). Now, compute the sum of the secondary diagonal (top right to bottom left), as efficiently as possible. (Hint: type `doc fliplr` in the Command window).

Problem 4.12.8:

Generate a set of 1,000 scores (`satscores`) in a 100×10 table that are normally distributed as ideal SAT scores (mean 500, $sd = 100$). Verify that the mean and standard deviation of each column are near 500 and 100, respectively.

Problem 4.12.9:

Make four 1×1000 matrices of random numbers, `v1`, `v2`, `v3`, and `v4`, using `randn`, and compute $W1 = v1 + v2$, $W2 = v1 + v3$, and $W3 = v3 + v4$. What are the correlations among `W1`, `W2`, and `W3`? If you've had a statistics course, explain the difference between the values of these correlations in terms of the sources of variability for `W1`, `W2`, and `W3`.

Problem 4.12.10:

Create *SATs*, a 1400×2 matrix of normally distributed random multiples of 10 (i.e. 200, 210, 220, etc.) with a mean of 500 and a standard deviation of 100, using `randn`, `round`, and other operations. Don't print it out! Find the mean and standard deviation of each column.

Problem 4.12.11:

The following code generates a 3×3 matrix, *A*, and reports the sum of the columns of *A*, as a row.

Code 4.12.11:

```
rng('default')
A = randi(9,3,3)
sum(A)
```

Output 4.12.11:

```
A =
     8     9     3
     9     6     5
     2     1     9
ans =
    19    16    17
```

Add exactly *one* character to the command `sum(A)`, to report the sum of the rows of *A*, as a row.

Add exactly *two* characters to the command `sum(A)`, to report the sum of the rows of *A*, as a column. (There are two solutions to this one!)

Add exactly *three* characters to the command `sum(A)`, to report the sum of all elements of *A*.

Add exactly *five* characters to the command `sum(A)`, to again report the sum of all elements of *A*.

Problem 4.12.12:

For the statistically minded, the matrix *A*, generated in Code 4.12.11 just after the random number generator was initiated, has three 9's, but no 4 or 7. Use MATLAB to compute the probability of the next (or any future) such matrix having *no* repeated digits.

Problem 4.12.13:

In Section 3.6 (Code 3.6.8), we noted that each value of a logarithmic series, after the first one, is a constant multiple of the prior value. Verify that this is the case (i.e., show the ratio between each element and the next) for `logseries1 = logspace(1,2,5)`, using shorthand notation for matrix indices (i.e., by writing one MATLAB statement).

5. Contingencies

This chapter covers the following topics:

- 5.1 Using the `if...else...end` construct
- 5.2 Using the `switch...case...end` construct
- 5.3 Using the `for...end` construct
- 5.4 Using the `while...end` construct and escaping from runaway loops
- 5.5 Vectorizing rather than using `for...end`
- 5.6 If-ing instantly
- 5.7 If-ing instantly once again and finding indices of satisfying values
- 5.8 Applying contingencies: Constrained random sequences and Latin squares
- 5.9 Practicing contingencies

The commands that are introduced and the sections in which they are premiered are:

<code>&</code>	(5.1)
<code><</code>	(5.1)
<code><=</code>	(5.1)
<code>==</code>	(5.1)
<code>></code>	(5.1)
<code>>=</code>	(5.1)
<code> </code>	(5.1)
<code>~=</code>	(5.1)
<code>ctrl-[</code>	(5.1)
<code>ctrl-]</code>	(5.1)
<code>ctrl-i</code>	(5.1)
<code>else</code>	(5.1)
<code>elseif</code>	(5.1)
<code>end (if)</code>	(5.1)
<code>exist (variable)</code>	(5.1)
<code>if</code>	(5.1)
<code>isempty</code>	(5.1)
<code>not</code>	(5.1)
<code>case</code>	(5.2)
<code>end (case)</code>	(5.2)
<code>otherwise</code>	(5.2)
<code>switch</code>	(5.2)
<code>end (for)</code>	(5.3)
<code>for</code>	(5.3)

<code>break</code>	(5.4)
<code>ctrl-c</code>	(5.4)
<code>end</code> (<i>while</i>)	(5.4)
<code>while</code>	(5.4)
<code>tic</code>	(5.5)
<code>toc</code>	(5.5)
<code>all</code>	(5.6)
<code>any</code>	(5.6)
<code>find</code>	(5.7)
<code>isnan</code>	(5.7)

5.1 Using the `if ... else ... end` Construct

The last chapter was concerned with calculations. The present chapter is concerned with contingencies. Contingencies are explicit rules for carrying out actions. A familiar example is going if a traffic light is green or stopping if it is red.

MATLAB has operators for directing traffic in much the same way. It uses Boolean operators to do so. Boolean operators yield one of two values: 1 (true) or 0 (false).

<code>==</code>	<code>a == b</code>	a equals b
<code>></code>	<code>a > b</code>	a is greater than b
<code>>=</code>	<code>a >= b</code>	a is greater than or equal to b
<code><</code>	<code>a < b</code>	a is less than b
<code><=</code>	<code>a <= b</code>	a is less than or equal to b
<code>~=</code>	<code>a ~= b</code>	a is not equal to b

Here is an example of how the first of these Boolean operators, `==` (the equals operator), is used. In this example, we employ an `if ... else ... end` construction. The program dictates that if `a` equals 1 (one condition), `b` should be multiplied by 2 (an action), or else (another condition), `b` should be multiplied by `-2` (another action).

Code 5.1.1:

```

b = 2;
a = 1;
if a == 1
    b = 2 * b;
else
    b = -2 * b;
end
b

```

Output 5.1.1:

```

b =
    4

```

Notice that a double equal sign is needed for the comparison `a == 1`. A single equal sign in an `if` statements would yield an error message. In MATLAB the double equal sign denotes the comparison for equality, whereas a single equal sign denotes assignment (`a = 1`, for example).

Code 5.1.1 also contains an `else` statement, indicating what to do if `a` does not equal 1. The program concludes with an `end` statement. This statement is mandatory for every `if` statement, whether or not there is an `else`. The `end` statement denotes completion of the range of code affected by the `if` statement.

Suppose you don't want to do anything if `a` does not equal 1. In that case, you can simply omit the `else` command as well as the command after it. Thus, `else` is optional.

Code 5.1.2:

```

b = 2;
a = -1;
if a == 1
    b = 2*b;
end
b

```

Output 5.1.2:

```

b =
    2

```

In the next example, `b` gets different values depending on whether `a` is negative, `a` equals 0, or `a` equals 1. The `elseif` command is useful in cases like this, where multiple comparisons are necessary. Trace through the code to verify that nothing will happen if `a` equals a positive non-zero value other than 1.

Code 5.1.3:

```
b = 2;  
a = 1;  
if a < 0  
    b = -1;  
elseif a == 0  
    b = 0;  
elseif a == 1  
    b = 1;  
end  
b
```

Output 5.1.3:

```
b =  
    1
```

Two Boolean operators can be combined by the & (and) and | (or) operators. We illustrate this principle by checking whether *a* is greater than or equal to 1 *and* less than or equal to 3. We must list each of these criteria formally, in contrast to the way we describe the criteria in everyday English, as in the previous sentence. We use `>=` to specify greater than or equal to, `&` to specify and, and `<=` to specify less than or equal to. (The parentheses in the `if` statement are optional but help clarify the intended parsing. We recommend using parentheses in contexts like this.) Note that the value of *b* is printed here only if it changes from its initial value.

Code 5.1.4:

```
b = 2;  
a = 2.7;  
if (a >= 1) & (a <= 3)  
    b = 2*b  
end
```

Output 5.1.4:

```
b =  
    4
```

In the next example, we check whether *a* is less than or equal to 1 or greater than or equal to 3. We use the `|` symbol to specify *or*. The value of *b* is unchanged if neither condition is met, but the value of *b* is written out in either case, whether it is changed or not.

Code 5.1.5:

```
b = 2;
a = 3.7;
if (a <= 1) | (a >= 3)
    b = -b;
end
b
```

Output 5.1.5:

```
b =
    -2
```

It is worth mentioning that the `|` symbol specifies just one kind of “or,” namely, “logical or.” When `|` is used in the above example, the condition is satisfied if *either* `a <= 1` *or* `a >= 3`. For information about other versions of `or` supported by MATLAB, type `help or` in the MATLAB command line.

In the next example we check whether the value of `a` differs from 10. We use the `~=` operator (“not equal to”) for this purpose.

Code 5.1.6:

```
b = -2;
a = 3.7;
if a ~= 10
    b = -b;
end
b
```

Output 5.1.6:

```
b =
     2
```

Nesting of `if` statements allows for more complex contingencies.

Code 5.1.7:

```
A = -2.3
a = 10
if A <= 0
    if a <= -5
        b = 1;    %if A is <=0 and if a <=-5, b gets 1
    else
        b = 2;    %if A is <=0 and if a is not <=-5, b gets 2
    end
end
```

```
else
    if a <= 5
        b = 3;    %if A is not <=0 and if a <=5, b gets 3
    else
        b = 4;    %if A is not <=0 and if a is not <=5, b gets 4
    end
end
b
```

Output 5.1.7:

```
A =
   -2.3000
a
    10
b =
     2
```

Notice that Code 5.1.7 uses indentation to accentuate the hierarchical nesting of the `for` and `if` statements. Using indentation greatly facilitates the analysis and debugging of code. Indentation occurs automatically as you type if you turn on “Apply smart indenting while typing” in the *language* section of MATLAB preferences. You can also select a block of text and indent it using `ctrl-]`. This keypress combination moves a selected block of text to the right to a previously defined tab position. Alternatively, you can outdent a selected block of text with `ctrl-[`. This moves the selected text to the *left*. Finally, if you have made some changes in an existing program, you can smart-indent after the fact by selecting the changed parts (or the entire program via `ctrl-a`) and hitting `ctrl-i`.

The `if` statement need not be used exclusively for computational decisions. It can also be used to delimit optional parts of your program so you can enable or disable them as needed. Often, you may want verbose output while developing a program, but when a program is fully reliable and in production, you may just want it to run as quickly as possible.

There are three ways to program `if` statements to control optional output. The quick and dirty way is to use the Boolean 1 (“true”) when you write the program and change the 1 to a 0 (“false”) when you don’t need the output any more:

Code 5.1.8:

```
if 1
    disp('Extensive output')
end
```

The quick and *clean* way to control a section of code, a bit more transparently, is to use `true` or `false` rather than 1 or 0 to control the `if` statement. Similarly, you can change the `true` to `false` when you don’t need the output any more, as in this example.

Code 5.1.9:

```

if false
    disp('Extensive output')
end

```

Finally, the *slow* and clean way to do the same thing is to use a Boolean at the beginning of your program to turn on and off *all* the verbose output in the program, at once:

Code 5.1.10:

```

verbose = true;
% ... other code of your program
if verbose
    disp('Extensive output')
end
% ... other code of your program
if verbose
    disp('More extensive output')
end

```

You can inhibit the verbose output by setting `verbose = false` at the beginning of the program. (There is nothing special about the word `verbose` here. We use it because the contingency concerns verbosity. However, we could have some other term, like `Give_Peace_A_Chance`.)

There are two other features of variables that you might need to check for in a program. The first is whether a particular variable exists. Here we check whether the variable *name* has been assigned, in which case the argument to `exist` is a string (in quotes) and the function returns 1 (true) if that name is present.

Code 5.1.11:

```

variable_1 = 5;
variable_1_exists = exist('variable_1')
variable_2_exists = exist('variable_2')

```

Output 5.1.11:

```

variable_1_exists =
    1
variable_2_exists =
    0

```

The second attribute of a variable that you might need to check is whether the variable is empty. Here we check whether the variable *value* is defined. We use `isempty`, assigning a variable (not a string in quotes) to the argument for `isempty`.

Code 5.1.12:

```
variable_3 = [];  
variable_4 = 9;  
variable_3_empty = isempty(variable_3)  
variable_4_empty = isempty(variable_4)
```

Output 5.1.12:

```
variable_3_empty =  
    1  
variable_4_empty =  
    0
```

5.2 Using the switch ... case ... end Construct

The `switch ... case ... end` construct is a convenient alternative to complex `if` statements. It compares a single variable against a number of possible values to execute alternative actions. Here is an example in which different actions are taken depending on the value of `year`. A variable representing a student's class year is used to determine what is printed. There is an optional catchall category, `otherwise`, for any cases that do not match one of the cases specified. For testing, we specify the class of 2017.

Code 5.2.1:

```
year = 2017    % For example  
switch year  
    case 2018  
        disp('First-year');  
    case 2017  
        disp('Sophomore');  
    case 2016  
        disp('Junior');  
    case 2015  
        disp('Senior');  
    otherwise  
        disp('Not a valid class year');  
end
```

Output 5.2.1:

```
year =  
    2017  
Sophomore
```

5.3 Using the `for ... end` Construct

The `for` loop lets you perform operations over and over, as many times as you specify. The variable that controls the `for` loop is available within the loop for computations. In this example, 2 is multiplied by the variable `i`, which takes on the values of 1, 2, 3, 4, 5, or 6 in successive passes through the loop. The `for` loop concludes with an `end` statement.

Code 5.3.1:

```
disp('      i      a');
disp(' ');
for i=1:6
    a=2*i;
    disp([i,a]);
end
```

Output 5.3.1:

i	a
1	2
2	4
3	6
4	8
5	10
6	12

In the next example, we add a semi-colon to the second line to suppress immediate output. In addition and more crucially, `i` also serves as the index for `a`. Thus, column 1 of `a` gets the product of 2×1 , column 2 of `a` gets the product of 2×2 , and so on.

Code 5.3.2:

```
for i=1:6
    a(i)=2*i;
end
a
```

Output 5.3.2:

a =	2	4	6	8	10	12
-----	---	---	---	---	----	----

It is also easy to use nested `for` loops to create more complicated matrices. In this example we set the element in the `i`-th row and `j`-th column of matrix `a` to `i+j` to make a matrix with six rows and three columns.

Code 5.3.3:

```
for i=1:6
    for j=1:3
        a(i,j)=i+j;
    end
end
a
```

Output 5.3.3:

```
a =
     2     3     4     8    10    12
     3     4     5     0     0     0
     4     5     6     0     0     0
     5     6     7     0     0     0
     6     7     8     0     0     0
     7     8     9     0     0     0
```

Wait a second! Hold on! Something very odd happened in Output 5.3.3! Code 5.3.3 specified a 6×3 matrix but we ended up with a 6×6 matrix. What happened?

The answer is that the variable `a` had not been cleared after it was used previously for a 1×6 matrix, so the values from Output 5.3.2 that were not overwritten were unchanged from Output 5.3.2 when Code 5.3.3 was run. We include this example as a reminder that MATLAB may incorporate new results into a previous matrix if that matrix is still active. To prevent this from happening (when it is not desired), it is advisable to clear the matrix that is being redefined or use a different variable name for each matrix.

Here is the same code as in the last example but with `a` cleared at the beginning.

Code 5.3.4:

```
clear a
for i=1:6
    for j=1:3
        a(i,j)=i+j;
    end
end
a
```

Output 5.3.4:

```
a =
     2     3     4
     3     4     5
     4     5     6
     5     6     7
     6     7     8
     7     8     9
```

When you create matrix indices using `for` loops, you must be sure to use positive integers. The following code produces an error.

Code 5.3.5:

```
for i=0:10
    a(i)=i+1;
end
```

Output 5.3.5:

```
??? Attempted to access a(0); index must be a positive
integer or logical.
```

The problem is that the first time a value was assigned to the i -th element of matrix a , i equaled 0, but a matrix can't have a zero-th element. The first element must have an index of 1, the second element must have an index of 2, and so. (You can't live in the 0-th house on a street...)

In case you conclude that negative numbers and 0 must be avoided in the context of `for` loops, consider this example, where i takes on values that are negative and, in one pass through the `for` loop, i is set to 0. Negative and zero values of i can be used in calculations. They just cannot be used as index values for arrays.

Code 5.3.6:

```
x=10;
disp('      i      a');
disp(' ')
for i=-3:3
    a=x*i;
    disp([i a]);
end
```

Output 5.3.6:

i	a
-3	-30
-2	-20
-1	-10
0	0
1	10
2	20
3	30

The following example shows that you can use `for` and `if` together. These elements are combined in the following program, where x is divided by i unless i is zero. (Remember, the result of dividing by zero is undefined).

Code 5.3.7:

```
x=10;
disp('      i      a')
disp(' ')
for i=-3:3
    if i~=0
        a=x/i;
        disp([i a]);
    end
end
```

Output 5.3.7:

```
      i      a
-3.0000  -3.3333
      -2      -5
      -1     -10
       1      10
       2       5
      3.0000   3.3333
```

What happens if you do not include the `if` statement in the last example (by commenting out the `if` and its corresponding `end`)?

Code 5.3.8:

```
x=10;
disp('      i      a')
disp(' ')
for i=-3:3
%     if i~=0
        a=x/i;
        disp([i a]);
%     end
end
```

Output 5.3.8:

```
      i      a
-3.0000  -3.3333
      -2      -5
      -1     -10
       0     Inf
       1      10
       2       5
      3.0000   3.3333
```

MATLAB forgives you for dividing by zero *but* alerts you to the misdemeanor by reporting the result as `inf` or by generating an error message (an option you can enable in the MATLAB preferences). Dividing by zero in some other programming languages causes the program to come to a grinding halt or, worse, the computer to crash. Despite MATLAB's forgiveness, it's wise not to divide by zero. Doing so may spoil your outputs (text or graphs) and will give you suspect results.

5.4 Using the `while ... end` Construct and Escaping from Runaway Loops

The `while ... end` construction lets you repeat an operation for as long as some condition holds. The `while ... end` construction is particularly helpful when it is difficult to anticipate how many steps will be needed until a condition changes state. Here is an example. The value of `a` is updated as long as `a` remains below 10.

Code 5.4.1:

```
a = 1;
b = .25;
steps = 0;
while a < 10
    a = a + a^b;
    steps = steps + 1;
end
a
steps
```

Output 5.4.1:

```
a =
    10.9475
steps =
     7
```

The `while` loop can be dangerous, however, because you can get caught in an endless `while` loop, as in the following program. This `while` loop was intended to work the same way as a `for` loop (`for x = 1:2:100`).

Code 5.4.2:

```
goal = 100;
x = 1;
while x ~= 100
    x^2
    x = x + 2;
end
```

Our hope was to square all odd integers below 100, starting with `x=1` and incrementing `x` in steps of 2 within the `while` loop. Because `x` was growing, the loop was expected to stop when `x` got big enough. The output was not what was expected, however, because `x` never had the value of exactly 100 that would permit escape from the `while` condition. As a result, MATLAB spewed forth a salvo of values that went well beyond 100. It continued growing until we pressed `ctrl-c` to stop it. Output 5.4.2 is not reproduced here because it was interrupted at an arbitrary point. Had we waited for the output to finish, we would never have gotten to this sentence!

It is not a good idea to get in the habit of relying on `ctrl-c` to escape from endless loops or from long listings of matrices caused by omissions of semi-colons. It is better to get in the habit of putting a semi-colon at the end of every line so outputs are suppressed. More important, however, is to think carefully and plan ahead to avoid runaway programs and other computationally unpleasant events.

One way to prevent endless loops when using `while` is to let the program run through only a limited number of steps, as in the following example. After 100 steps, the `break` command is executed. When the `break` command is invoked, the program breaks out of the `while` loop containing it. The `break` command can also be used to break out of `for` loops when some condition indicates that the loop should be prematurely terminated. In the output below, we omit intermediate material that the computer actually produced.

Code 5.4.3:

```
goal = 100;
x = 1;
steps = 1
while x ~= 100
    x^2
    x = x + 2;
    steps = steps + 1;
    if steps > 100
        break
    end
end
```

Output 5.4.3:

```
steps =
     1
ans =
     1
ans =
     9
ans =
    25
ans =
    49
```

```

ans =
    81
% ... output omitted
ans =
    38025
ans =
    38809
ans =
    39601

```

Another way to avoid endless loops is to recognize that exact comparisons may be never met (as in Code 5.4.2). When you suspect this might happen, use a different comparison. If the fourth line had been `while x <= 100` instead of `while x ~= 100`, the program would have worked as intended, stopping when `x` exceeded 100.

Notwithstanding all of the cautionary remarks given above, sometimes *potentially* endless loops can be useful. A potentially endless loop begins with `while true` or, equivalently, `while 1`, and then is escaped when some condition is met. Consider trying to generate a special random sequence to describe three experimental conditions (represented by 1, 2, and 3) to be run repetitively. The numbers 1, 2, and 3 are supposed to appear three times in random order, but you don't want any consecutive repetitions of the same number. Thus, `[1 3 2 3 2 1 2 3 1]` would be OK, but `[1 3 2 2 3 1 2 3 1]` would not be, because of the repeated '2' in positions 3 and 4. You could find an acceptable sequence by just trying random sequences over and over until you get one that fits the conditions, at which time you could use the `break` command to get out of the `while` loop.

Code 5.4.4:

```

clc
outputForTesting = true;
while true
    candidate = [randperm(3) randperm(3) randperm(3)];
    if candidate(3) ~= candidate(4) & ...
        candidate(6) ~= candidate(7);
        break
    end
    if outputForTesting
        badcandidate = candidate
    end
end
goodsequence = candidate

```

Output 5.4.4:

```

badcandidate =
    1     2     3     3     2     1     3     1     2
badcandidate =
    3     2     1     3     1     2     2     3     1

```

badcandidate =	3	2	1	1	3	2	1	3	2
badcandidate =	3	2	1	3	2	1	1	2	3
goodsequence =	2	3	1	2	3	1	2	1	3

In practice, this approach might be an efficient way to generate constrained sequences that would be difficult to generate otherwise. For example, when Code 5.4.4 was tested, the first or second candidate was usually acceptable, and the longest sequence of bad candidates was 7. In the next run, four bad outputs came along before a good one materialized. (Once you know the program works, you could change `outputForTesting = true;` to `outputForTesting = false;` to suppress the no-longer-needed output.)

5.5 Vectorizing Rather than Using `for ... end`

Earlier in this chapter you were introduced to `for` loops. These are useful when `if` statements are nested within them (as in Code 5.3.7 and Code 5.3.8), or when other `for` loops are nested within them (as in Code 5.3.3 and Code 5.3.4). However, `for` loops run slowly relative to instructions that can be completed, from the programmer's point of view, in one fell swoop.

One way of increasing computational efficiency is to avoid `for` loops by exploiting MATLAB's vector capabilities. The term used in the MATLAB programming community for giving such all-in-one instructions is *vectorizing*. When instructions are vectorized, processing time can be greatly reduced.

You have already been exposed to vectorizing in this book, though you didn't see the term before. In Chapters 3 and 4 (before the `for` loop was introduced), you saw how values were assigned to matrices in single statements. For example, in Code 3.6.2, the values `[1:6]` were assigned to a matrix `M` simply by writing `M=[1:6]`. This is an example of vectorizing. It turns out that it takes less time to assign the values 1 through 6 to the first six elements of `M` by vectorizing than by using a `for` loop and saying, via code, "if the current index is 1, then `M(1)` gets 1, if the current index is 2, then `M(2)` gets 2, and so on." Similarly, you learned how to make a matrix of numbers using, for example, `zeros(100)`, `ones(50,200)`, or `randi(8,100,100)`. Each of these operations could be accomplished with one or more nested `for` loops, but it is faster to do the operations directly. Clearly, for very small matrices, the time difference is immaterial, but for larger matrices or more complex calculations, vectorizing can make a noticeable difference.

Why does time grow appreciably when loops are used? Consider a `for` loop that goes through many cycles. In each cycle some processing resources (time) must be devoted to starting the loop and testing if it needs to be repeated. In addition, when a matrix is first used, MATLAB is parsimonious in the amount of memory assigned for it (roughly 1,000 elements). If a matrix grows within the loop, however, MATLAB must interrupt its computations to allocate more memory for the matrix, and this process takes time.

Consider the example of the pickup truck from Section 3.8. It would be very inefficient, if you were moving, to first rent a pickup truck, and then, when it was full, return it to the rental agency and rent a box truck, and then, when *that* got full, return it to rent a moving van, and so on. Better to rent a truck as large or somewhat larger than you need right from

the start. In MATLAB, renting an adequate truck is accomplished by pre-allocation, that is, by initially generating a large enough matrix variable to accommodate its eventual size.

Here is an example that shows how much more slowly `for` loops can take than vectorizing. The program that achieves the demonstration measures its computation time using a handy stopwatch function provided by the MATLAB commands `tic` and `toc`. As its name suggests, `tic` starts a stopwatch (by reading the computer's clock) and `toc` reports the stopwatch value (by reading the computer clock and computing the time elapsed since the last `tic`). Elapsed time is reported in seconds and fractions of a second.

Code 5.5.1 uses `tic...toc` to show that it may take more time to assign values to a matrix with a `for` loop than it does to assign the same values to a matrix by vectorizing. The program also illustrates that some time can be saved by pre-allocating memory.

There are several parts of the program. The first computes the time for defining a million random numbers directly (using the `randn` command to make a 1000×1000 matrix). The second part generates a 1000×1000 matrix and fills it with random numbers one at a time, within two nested `for` loops. The third part does the same, but first pre-allocates memory by generating a 1000×1000 matrix of all zeros. (The execution times reported are not fixed. It will depend on your hardware and the state of your machine when the program is run).

Code 5.5.1:

```
% Part 1: Generate numbers using RANDN
clc
close all
clear
tic
r = randn(1000,1000);
SecondsToGenerateMillionRandom_Directly = toc
% Part 2: Generate numbers using FOR, without preallocation
clear r
tic
for ii = 1:1000
    for jj = 1:1000
        r(ii,jj) = randn(1,1);
    end
    t(ii) = toc;
end
SecondsToGenerateMillionRandom_Forloops = toc
% Part 2: Generate numbers using FOR, with preallocation
clear r
tic
r = zeros(1000,1000);
for ii = 1:1000
    for jj = 1:1000
        r(ii,jj) = randn(1,1);
        tpa(ii) = toc;
    end
end
end
```

```
SecondsToGenerateMillionRandom_Preallocated_
ThenForLoops = toc
```

Output 5.5.1:

```
SecondsToGenerateMillionRandom_Directly =
    0.0502
SecondsToGenerateMillionRandom_Forloops =
    4.0844
SecondsToGenerateMillionRandom_Preallocated_
ThenForLoops =
    2.3898
```

Using the most direct method, `randn` generated a million random numbers in just 50 ms. When `for` loops were used, it took more time (more than 2 seconds), and it took even longer (4 seconds) if the matrix had not been not pre-allocated and had to grow within the `for` loop. The fastest way to generate a matrix, then, is to do so directly by vectorizing.

The point of this example is not to discourage you from using `for` loops in all cases, because `for` loops can sometimes be easier to understand than vectorizing, especially if you are relatively new to MATLAB. Furthermore, in some cases, `for` loops are essential (e.g., when `if` statements or other `for` loops are nested within them).

As you gain more expertise with MATLAB, keep in mind that `for` loops should be used judiciously. In addition, it is a good idea to pre-allocate memory for potentially large matrices, when the matrices risk growing within a loop, or if you need to have precise control over the computation time of your program.

5.6 If-ing Instantly

Just as assignments can be achieved without one-at-a-time instructions, comparisons can be achieved in single statements. We call this *if-ing instantly*, our shorthand (not an official MATLAB term) for testing an array of values simultaneously. Here's an example of if-ing instantly.

Code 5.6.1:

```
a = [12 15 17 13 15 12 14];
b = ( a > 13 )
```

Output 5.6.1:

```
b =
    0     1     1     0     1     0     1
```

The `b` array consists of only 1's and 0's. Why? They result from the Boolean operator `a > 13` (see Section 5.1). So `b` is an array of truth values where, by convention, 1 means

“true” and 0 means “false.” The second, third, fifth, and seventh elements of `b` are true because the second, third, fifth, and seventh elements of `a` met the condition of being greater than 13.

We can use this array of truth values to go back and find all the elements of `a` that satisfy the test. Conveniently, MATLAB has a useful function called `find` to do that kind of task. Here we use `find` to identify the true elements of `b`.

Code 5.6.2:

```
a = [12 15 17 13 15 12 14];
b = ( a > 13 );
indices_of_good_values_in_a = find(b)
the_good_values_themselves = ...
    a(indices_of_good_values_in_a)
```

Output 5.6.2:

```
indices_of_good_values_in_a =
     2     3     5     7
the_good_values_themselves =
    15    17    15    14
```

Take your time “breathing in” this code. The first line returns the indices or the addresses (in this case, the column numbers) of `b` that are true; the second, third, fifth, and seventh elements of `b` are greater than 13. The second line of Code 5.6.2 returns the values of `a` with those “true indices.”

It turns out that there is an even more direct shortcut for if-ing instantly that does not require the second variable (`b` in Code 5.6.2). This is done using the matrix returned by the Boolean expression as the index for the matrix being tested.

Code 5.6.3:

```
a = [12 15 17 13 15 12 14];
the_good_values_instantly = a( a>13 )
```

Output 5.6.3:

```
the_good_values_instantly =
    15    17    15    14
```

What’s happening here is that `a( a>13 )` is *immediately* finding all the values of `a` that meet the specified condition. (When the first author uses constructions like `a( a>13 )`, he says to himself “a such that a is greater than 13.”)

You can take the same general approach to comparing two matrices. Suppose you want to find the values in one matrix that are the same as in another matrix. In the following, `find` returns the indices of elements meeting the criterion.

Code 5.6.4:

```

m1 = [
    16    13     3     2
     9    12     6     7
     5     8    10    11
     4     1    15    14];
m2 = [
    16     2     3    13
     5    11    10     8
     9     7     6    12
     4    14    15     1];
cells_inwhich_m1_equals_m2 = (m1 == m2)
indices_of_the_equal_values = find(m1 == m2)
the_values_that_are_equal = m1(m1 == m2)

not_m1_equals_m2 = ~(m1 == m2)
m1_notequal_m2 = (m1 ~= m2)

```

Output 5.6.4:

```

cells_inwhich_m1m1_equals_m2 =
    1     0     1     0
    0     0     0     0
    0     0     0     0
    1     0     1     0
indices_of_the_equal_values =
    1
    4
    9
   12
the_values_that_are_equal =
   16
    4
    3
   15

not_m1_equals_m2 =
    0     1     0     1
    1     1     1     1
    1     1     1     1
    0     1     0     1
m1_notequal_m2 =
    0     1     0     1
    1     1     1     1
    1     1     1     1
    0     1     0     1

```

Sometimes you want to know if *any*, or *every*, element of a matrix meets some condition. The functions `any` and `all` serve that purpose.

Code 5.6.5:

```

mymatrix = magic(3)
any_5_bycolumns = any(mymatrix == 5)
all_lessthan_or_equal_8_bycolumns = all(mymatrix <= 8)
any_5_inthe_wholematrix = any(mymatrix(:) == 5)
all_lessthan8_inthe_wholematrix = all(mymatrix(:) < 8)

```

Output 5.6.5:

```

mymatrix =
     8     1     6
     3     5     7
     4     9     2
any_5_bycolumns =
     0     1     0
all_lessthan_or_equal_8_bycolumns =
     1     0     1
any_5_inthe_wholematrix =
     1
all_lessthan8_inthe_wholematrix =
     0

```

5.7 If-ing Instantly Once Again and Finding Indices of Satisfying Values

The code below provides another example of instant if-ing that illustrates another type of construction, one that relies on the “or” operator (`|`). Here the matrix `h` is assigned the numbers 1 through 11, which are randomly permuted and added to 10 to generate a random series ranging from 11 to 21. The third line of Code 5.7.1 uses `find` to return the indices of elements of `h` that are equal either to 12 or to 16. The `find` function can be useful for determining which participants satisfied one or more conditions in an experiment.

Code 5.7.1:

```

h=randperm(11)+10
h==12|h==16
find(h==12|h==16)
values_sought = h(find(h==12|h==16))

```

Output 5.7.1:

```

h =
    16    11    17    14    21    19    15    18    13
    20    12
ans =
     1     0     0     0     0     0     0     0     0
     0     1

```

```
| ans =  
|      1      11  
| values_sought =  
|      16      12
```

If-ing instantly has many applications. One of its virtues is that it is executed quickly. Another is that it is simpler to program than a `for` loop, and so is more likely to be right the first time. You saw an example of if-ing instantly before in this book, in Section 4.6.2, in the expression `Data(not(isnan(X)))`. There, we wanted to know which elements of an array `Data` were not NaN's, so we tested the entire array with one expression. In general, using a logical expression in the place of a matrix's index is an extremely useful shortcut for selecting subsets of a matrix. Here is an example, where for a set of random integers, you wish to set all values greater than 50 to 50.

Code 5.7.2:

```
| z = randi(100,1,8)  
| z(z > 50) = 50
```

Output 5.7.2:

```
| z =  
|      82      91      13      92      64      10      28      55  
| z =  
|      50      50      13      50      50      10      28      50
```

The values 82, 91, 92, 64, and 55 have all been trimmed to 50 with one command.

Similarly, if you wanted to count the number of elements that are less than 25, you could do so as follows.

Code 5.7.3:

```
| length(z(z < 25))
```

Output 5.7.3:

```
| ans =  
|      2
```

A logical expression that returns values meeting some condition can even control a `for` loop if there is an operation to be performed on only those values of a variable that meets a condition. This example reports all values evenly divisible by 3.

Code 5.7.4:

```
| A = [ 1 4 6 3 8 6 5 9 2 7 5];  
| disp('The following elements of A are divisible by 3:')
```

```

for i = A(mod(A,3)==0)
    disp(i)
end

```

Output 5.7.4:

```

The following elements of A are divisible by 3:
    6
    3
    6
    9

```

5.8 Applying Contingencies: Constrained Random Sequences and Latin Squares

Now that you have learned about many of the basics of MATLAB programming, you can begin to explore their applications in behavioral research. The examples below are based on practical situations that the authors have encountered in their research projects. For all of the applications in this and subsequent chapters, before reading the code we have generated, we invite you to give some thought to how *you* might solve the problem.

The first application is an elaboration on the generation of constrained random sequences from Code 5.4.4. The challenge is to present 12 successive stimuli so there are an equal number of stimuli on the left and right, but no more than three consecutive stimuli on either side. In our approach, we use '1' to represent the left side, and '2' to represent the right. The first part of the program sets up an initial array, `trialsequence`, consisting of an equal number of left and right trials, and sets the Boolean `done` to `false`. The program will end when `done` is `true`.

In the second part of the program, we begin by (optimistically) setting `done = true`, before testing for the condition. The program then randomizes the order of the sides using `randperm` and tests the new sequence against the criterion. Because we are randomizing a sequence that already has five stimuli on each side, we don't have to worry about that constraint any more. However, any sequence of four identical trials would violate the constraint that there be no more than three consecutive stimuli on the same side, and if any such are found, we set `done = false`. The program finishes if `done` is still `true` at the end of the `while` loop, and then reports the valid sequence. It also reports how many tries it took, which we count in the variable `cycles`. This lets us confirm that we can find a valid sequence without too many tries. In 20 tests of the program, the largest number of cycles observed was 6. The modal value was, reassuringly, 1.

Before you rush to adopt this algorithm for generating other sequence lengths, heed this cautionary note. Imagine you wanted sequences of 1,000 trials (500 left and 500 right) such that there were no runs of three on one side. It might take a very long time to hit on a satisfactory sequence by chance, as in Code 5.8.1. In such a case it would be preferable to find another method for generating the sequence, such as generating subsets of the sequence for shorter blocks of trials.

Code 5.8.1:

```

rng('shuffle')
clc
sequenceOf1sAnd2s = [ones(1,6) ones(1,6)*2];
done = false;
cycles = 0;

while not(done)
    cycles = cycles + 1;
    done = true;
    sequenceOf1sAnd2s = sequenceOf1sAnd2s(randperm(12));
    for i = 4:12
        % Detect any runs of 1's or 2's
        if sequenceOf1sAnd2s(i) == ...
            sequenceOf1sAnd2s(i-1)...
            & sequenceOf1sAnd2s(i) == ...
            sequenceOf1sAnd2s(i-2)...
            & sequenceOf1sAnd2s(i) == ...
            sequenceOf1sAnd2s(i-3)
            done = false;
        end
    end
end

sequenceOf1sAnd2s
cycles

```

Output 5.8.1:

```

sequenceOf1sAnd2s =
     2     1     2     1     2     2     2     1     1
1     2     1
cycles =
     2

```

A third example that exploits what you have just learned about randomization is the generation of Latin squares. A Latin square is an $n \times n$ matrix with the defining properties that each of the integers 1 through n occurs exactly once per row and exactly once per column. Latin squares are often used in experimental designs to ensure, for example, that each condition (represented by 1 through n) is experienced once by each subject (represented by a row) and that, for the experiment as a whole, each condition occurs once in each position of the experiment series (represented by a column).

To generate a Latin square, you can begin by generating a non-random matrix that has each integer once and only once in each row and column. This can be done by putting the integers $1 \dots n$ in the first row, and then making each successive row be a shift (by one element) of the preceding row. A unique Latin square can then be generated by first randomizing the order of the rows (which does not change the matrix's defining properties) and then randomizing the order of the columns (which also preserves the defining properties).

Code 5.8.2:

```

rng('shuffle')
clc
clear m
LSSize = 7;
% The first line of m is the integers 1:LSSize
m(1,:) = [1:LSSize];

% Each subsequent line is the preceding line, rotated by
% one element
for i = 2:LSSize
    m(i,:) = m(i-1,[2:end,1]);
end
OriginalMatrix = m
% Permute rows, ...
m1 = m(randperm(LSSize),:);
RowsPermutedMatrix = m1
% ...then permute columns to make randomized Latin
% Square matrix
m2 = m1(:,randperm(LSSize));
LatinSquareMatrix = m2

```

Output 5.8.2:

```

OriginalMatrix =
     1     2     3     4     5     6     7
     2     3     4     5     6     7     1
     3     4     5     6     7     1     2
     4     5     6     7     1     2     3
     5     6     7     1     2     3     4
     6     7     1     2     3     4     5
     7     1     2     3     4     5     6

RowsPermutedMatrix =
     1     2     3     4     5     6     7
     5     6     7     1     2     3     4
     7     1     2     3     4     5     6
     4     5     6     7     1     2     3
     2     3     4     5     6     7     1
     3     4     5     6     7     1     2
     6     7     1     2     3     4     5

LatinSquareMatrix =
     2     6     3     7     4     5     1
     6     3     7     4     1     2     5
     1     5     2     6     3     4     7
     5     2     6     3     7     1     4
     3     7     4     1     5     6     2
     4     1     5     2     6     7     3
     7     4     1     5     2     3     6

```

5.9 Practicing Contingencies

Try your hand at the following exercises, using only the methods introduced so far in this book or in information given in the problems themselves.

Problem 5.9.1:

You want to show stimuli to a participant in a psychophysics experiment. The stimuli to be shown should have values of A^B , where A takes on the values of 1, 2, 3, and 4, and B takes on values of 1, 2, 3, and 4. Write a program to generate the 16 stimulus values.

Now reorder the values into a random presentation order.

Problem 5.9.2:

The following matrix contains fictional data from a reaction-time experiment. Each row contains the mean reaction time and proportion correct for a different participant. Use a `for` loop and an `if` statement to identify the participants who had mean reaction times greater than 500 ms and proportions correct greater than .65. The output should contain two matrices, called `Identified_Participants` and `OK_Scores`. The values in `Identified_Participants` should be the numbers of the participants fulfilling the criteria. The values in `OK_Scores` should be rows, each with two columns, one for reaction time and one for proportion correct.

```
RT_and_PC_Data = [  
    390  .45  
    347  .32  
    866  .98  
    549  .67  
    589  .72  
    641  .50  
    777  .77  
    702  .68  
];
```

Problem 5.9.3:

Use logical comparisons instead of a `for` loop and an `if` statement to solve the last problem.

Problem 5.9.4:

Find out how long it takes your computer to identify values greater than the overall mean in a 1000×1000 matrix of random numbers, using `for` and `if` statements. Also find out how long it takes your computer to identify values greater than the overall mean through instant `if`-ing. Because the matrix is large, you will want to suppress most other output.

Problem 5.9.5:

You are curious to know how many trials it takes a participant to get a requisite number of responses correct (trials to criterion) in a categorization task. You are especially interested to know how the trials to criterion depend on the participant's learning rate. Suppose there are four category names whose corresponding stimuli are presented equally often. Suppose as well that participants are told the correct response after each response. Suppose finally that the probability of a correct response is a logarithmic function of the number of completed trials, according to the equation:

$$p_correct = base_rate + learning_rate * \log(trial),$$

where `trial` can take on the values 1, 2, 3, ..., 200, `learning_rate` can be any real number between 0 and 1, `base_rate` equals 1/4 (i.e., 1 over the number of categories), and `p_correct` cannot exceed 1. Write a program that lets you explore the effects of `learning_rate` on number of trials to criterion. You can set the criterion to whatever value(s) you choose.

Problem 5.9.6:

Solve for the root(s) of the quadratic equation $ax^2 + bx + c = 0$. Use appropriate contingencies (`if` or `switch` statements) to report the different cases (two roots, one root, and no roots) depending on the values of the coefficients a , b , and c , and use `disp` to describe the results you report. By way of reminder, a quadratic equation has two roots, x_1 and x_2 , unless $b^2 = 4ac$, in which case there is just one root, and unless $b^2 < 4ac$, in which case there are no roots (no real values of x that solve the equation).

$$x_1 = \frac{-b + \sqrt{b^2 - 4ac}}{2a} \quad \text{and} \quad x_2 = \frac{-b - \sqrt{b^2 - 4ac}}{2a}$$

Try your program with at least these three sets of values:

```
a = 16; b = 0; c = -4;
a = 9;  b = 6; c = 1;
a = 9;  b = 0; c = 1;
```

Problem 5.9.7:

Apply the trimming technique of Code 5.7.4 to the output of Problem 4.12.8 to limit the range of SAT scores to $200 \leq \text{SAT} \leq 800$. Count the number of 200's and 800's in the data set after you have done so.

Problem 5.9.8:

Write a program to compute the standard error of the mean of a uniform distribution (use `rand`) that has n values (a value you specify for each run of the program). Build in a

contingency to divide the standard deviation by the square root of n if n is greater than or equal to 30, but to divide by the square root of $n-1$ if n is less than 30.

Problem 5.9.9:

Create a 2×100 matrix whose first and second rows are the numbers 1 to 100. Then multiply columns 3, 6, 9, 12, and 15 by 3.

Problem 5.9.10:

Having solved problem 5.9.9, see if you can achieve the same thing by typing the numbers 1 and 100 only once and by never typing the numbers 6, 9, or 12. Hint: Who ever said an index can only be one number?

Problem 5.9.11:

The following code will generate a set of student data in which column 1 is the student number, column 2 is an integer representing the class year (2015, 2016, 2017, or 2018), and column 3 is the student's grade point average (which ranges from 2.0 to 4.0). Since the random number generator is initialized at the beginning, you will get the same sequence as we did, and your checkvalues should agree with the values reported in the comment lines.

```
clc
clear
rng('default')
data(:,1) = randperm(300);
data(:,2) = randi(4,300,1) + 14;
data(:,3) = randi(20,300,1)/10 + 2;
checkvalues = mean(data)
% checkvalues should be:
%   150.5000   16.4533    2.9837
```

Without printing out the data matrix, answer the following: What students (by student number) had a 4.0 average? Who are the seniors (class of '15) who will graduate with honors (GPA ≥ 3.5)? How many first-year students (class of '18) are likely to elect to be psychology majors, as predicted by their GPA being greater than 3.0? What is the GPA of student #1 (the student with that student number, not necessarily the first student in the matrix)? What is the standard deviation of the GPAs of second-year students (class of '17)?

Problem 5.9.12:

To make “truly random” numbers available to the scientific community, some years ago the RAND corporation published a list of a million random digits (RAND Corporation, 1955). (The volume is still available through Amazon's “print on demand” service. An Amazon reviewer self-identified as *a curious reader* observed, “Such a terrific reference work! But

with so many terrific random digits, it's a shame they didn't sort them, to make it easier to find the one you're looking for.") Using `tic`, `toc`, and `randi`, generate a million random digits in a $100 \times 100 \times 100$ array, measuring how long it takes to generate the digits in three ways: using three nested `for` loops without pre-allocation of the $100 \times 100 \times 100$ array; using three nested `for` loops with pre-allocation; and directly (be sure to clear the array at the beginning of each generation). Also compare the times for setting a 1×1000000 array to the value `NaN`, directly and using `for` loops with and without pre-allocation. (If this seems to take *for-ever* on your particular machine, interrupt your program with `ctrl-c` and try again with a smaller number.)

Problem 5.9.13:

Adam, Beth, Charlie, and Deb share an apartment. The dishes need to be washed every evening, and the residents agree to follow a rotating schedule, starting with Adam on the first evening. Write a program to print the date and the responsible resident for a particular day of the month. Who is to do the dishes on the 13th day? Begin your program with `today = 13`, but make it so it would work if that first line specified any day of the month up to 31.

Since no months except February have a multiple of four days, Adam would get the short end of this deal if he were the first to go each month. How might you address this dilemma, other than negotiating a small compensatory rent reduction for Adam? Hint: Doing so would require you to modify just one variable value in your working program at the beginning of each month.

After you have the program working, test your program for all possible days of the month, by omitting the first line (`today = 13`) and wrapping a `for` loop around your program:

```
for today = 1:31

    % your program goes here ...

end;
```

6. Input-Output

This chapter covers the following topics:

- 6.1 Copying and pasting data by hand
- 6.2 Getting input from a user and displaying the result
- 6.3 Pausing
- 6.4 Recording reaction time and other delays with `tic ... toc`
- 6.5 Formatting numbers for screen outputs
- 6.6 Assigning arrays of literal characters (strings) to variables
- 6.7 Controlling file print formats
- 6.8 Writing data to named files
- 6.9 Writing text to named files
- 6.10 Checking and changing the current directory
- 6.11 Reading data saved as plain text from named files
- 6.12 Reading data from and writing data to Excel spreadsheets
- 6.13 Taking precautions against overwriting files
- 6.14 Saving and loading variables in native MATLAB format
- 6.15 Learning more about input and output
- 6.16 Practicing input-output

The commands that are introduced and the sections in which they are premiered are:

<code>' '</code> (<i>string delimiter</i>)	(6.2)
<code>input</code>	(6.2)
<code>pause</code>	(6.3)
<code>format</code>	(6.5)
<code>format bank</code>	(6.5)
<code>format compact</code>	(6.5)
<code>format long</code>	(6.5)
<code>format long g</code>	(6.5)
<code>format loose</code>	(6.5)
<code>format rat</code>	(6.5)
<code>format short</code>	(6.5)
<code>format short g</code>	(6.5)
<code>' '</code> (<i>apostrophe in string</i>)	(6.6)
<code>'s'</code>	(6.6)
<code>\n</code>	(6.6)
<code>\t</code>	(6.6)
<code>%%</code>	(6.6)
<code>%d</code>	(6.6)
<code>%e</code>	(6.6)

<code>%f</code>	(6.6)
<code>%s</code>	(6.6)
<code>sprintf</code>	(6.6)
<code>fprintf</code>	(6.7)
<code>dlmwrite</code>	(6.8)
<code>fclose</code>	(6.8)
<code>fopen</code>	(6.8)
<code>type</code>	(6.8)
<code>cd</code>	(6.10)
<code>dir</code>	(6.10)
<code>ls</code>	(6.10)
<code>pwd</code>	(6.10)
<code>load (.txt file)</code>	(6.11)
<code>xlsread</code>	(6.12)
<code>xlswrite</code>	(6.12)
<code>clock</code>	(6.13)
<code>exist</code>	(6.13)
<code>load (.mat file)</code>	(6.14)
<code>save (.mat file)</code>	(6.14)
<code>fgetl</code>	(6.15)
<code>fread</code>	(6.15)
<code>fseek</code>	(6.15)
<code>fwrite</code>	(6.15)
<code>iofun</code>	(6.15)
<code>textscan</code>	(6.15)

6.1 Copying and Pasting Data by Hand

In all of the examples presented so far, matrices have been generated with little control of their format, either for input or for output. In addition, matrices have been output only to the MATLAB Command window. It would be desirable to have more control of input and output, especially for large data sets. This chapter covers ways of doing this.

One way of getting data into a program is to copy them from another source, such as Microsoft Word or Excel. A method that can be used for this purpose is to create an assignment command in the Editor window, leaving space between the opening and closing brackets of the matrix that you want to import from Word or Excel and then pasting text between those brackets. Here is an example of code that can be created prior to pasting.

Code 6.1.1:

```
| my_data = [  
|           ]
```

Output 6.1.1:

```
| my_data =  
|      []
```

Having written this code, you can paste text into it. In this case, a 2×4 matrix is pasted in, consisting of the numbers 1 through 4 in the first row and the numbers 5 through 8 in the second.

Code 6.1.2:

```
| my_data = [  
| 1 2 3 4  
| 5 6 7 8  
|           ]
```

Output 6.1.2:

```
| my_data =  
|      1      2      3      4  
|      5      6      7      8
```

One reason for considering this example is to show that the closing bracket for a matrix does not have to be on the same line as the opening bracket, though the opening bracket does have to be on the same line as the equal sign (=). Another reason for considering this example is to mention that it is safer to paste simple text than formatted text into .m files. For example, copying several cells of data from an Excel spread sheet can yield unexpected results. Similarly, pasting nicely formatted data from Word may make a mess once it's in MATLAB. If you're going to paste data into a .m file, first convert the data to plain text.

Copying and pasting can also be used to transfer the output of a MATLAB program to another file, such as a Word document. If you are planning to do this—later in this chapter we will show you more direct ways of getting data and text into and out of MATLAB—then, after generating a matrix with MATLAB, you can select the output from MATLAB's Command window, and copy and paste it into the other document.

6.2 Getting Input from a User and Displaying the Result

How else can data be entered into MATLAB? One context in which this question can be addressed is a situation commonly encountered in behavioral science: gathering data interactively. Suppose you want someone to input data to a computer. The challenge is to design an interactive mode of communication that ensures that the data come in both as you wish and as the user wishes (provided the user is being cooperative).

A function that is useful in this context is `input`. The `input` function takes as its first argument a prompt string. In the example that follows, the prompt is, “What is your favorite number?” When MATLAB encounters the `input` command, it displays the string provided as the argument. Notice that apostrophes (`' '`) surround the text to mark it as a *string*, a matrix of alphabetic characters rather than numbers. Putting a space between the question mark and the final apostrophe leaves a space between the question mark and the user’s typed response. The output appears in the Command window.

Code 6.2.1:

```
| favorite = input('What is your favorite number? ')
```

Output 6.2.1a:

```
| What is your favorite number?
```

MATLAB waits for a number to be typed in, and next waits for the <enter> or <return> key to be pressed. If the user types “3,” here is what appears in the Command window.

Output 6.2.1b:

```
| What is your favorite number? 3
| favorite =
|      3
```

When using `input`, it is important to “idiot-proof” the interaction. The term “idiot-proof” conveys the idea that users—even well-intentioned, perfectly intelligent ones—may sometimes do unexpected things, such as hitting keys that generate bad data. Consider the following exchange.

Code 6.2.2:

```
| favorite = ...
|      input('What is your favorite number between 2 and 7?')
```

If the user accidentally types an alphabetic character such as `p` rather than a number, but `p` is not a defined variable, MATLAB sends an error message because only a valid MATLAB expression (such as a number) is acceptable in this context.

Output 6.2.2:

```
| ??? Undefined function or variable 'p'.
```

Even if the user types a number, there is no guarantee it will be useful. For example, if the user types a number outside the range 2 to 7, you are stuck with that value, which may be inconvenient later.

A strategy for idiot-proofing the interaction is to exploit the `while ... end` loop (see Chapter 5). In the following example, the user is asked for his or her favorite number as long as the value of `favorite` is less than 2 or greater than 7. A `while` loop is used for this purpose. To make sure the `while` loop is entered, `favorite` is initialized to a value less than 2 or greater than 7. A convenient initialization value is `-inf`.

Code 6.2.3:

```
favorite = -inf;
while (favorite < 2) | (favorite > 7)
    favorite = ...
    input('What is your favorite number between 2 and 7? ')
end
disp('OK, got it!')
```

Output 6.2.3:

```
What is your favorite number between 2 and 7? 88
favorite =
    88
What is your favorite number between 2 and 7? 0
favorite =
    0
What is your favorite number between 2 and 7? 3
favorite =
    3
OK, got it!
```

As shown here, the user eventually figures out that there is a problem with his or her answer. However, not all users are as patient or as diligent as one hopes. Consequently, it may help to provide more polite or informative prompts, as illustrated below, where `disp` is used to display a warning message as well as a final congratulatory message.

Code 6.2.4:

```
favorite = 0;
while (favorite < 2) | (favorite > 7)
    favorite = ...
    input('What is your favorite number between 2 and 7? ')
    if (favorite < 2) | (favorite > 7)
        disp('Sorry, not a valid number between 2 and 7.')
        disp('Try again, please.')
    end
end
disp('OK, got it!')
```

Output 6.2.4:

```

What is your favorite number between 2 and 7? 1
favorite =
    1
Sorry, not a valid number between 2 and 7.
Try again, please.
What is your favorite number between 2 and 7? 8
favorite =
    8
Sorry, not a valid number between 2 and 7.
Try again, please.
What is your favorite number between 2 and 7? 5
favorite =
    5
Got it!

```

6.3 Pausing

Sometimes you can help users feel a little less harried by slowing things down. The pause command is handy for this purpose. The following code shows how the pause command is used. The program first uses `disp` to show the message to which the user should respond. Then the computer is told to pause until a key (any key) is struck. For the key to take effect, the Command window must be active, so make sure to activate the Command window, using the `commandwindow` instruction before (though not necessarily immediately before) the `pause` command is issued. In the code below, the Command window is activated right before the `pause` command is given.

Code 6.3.1:

```

disp('Hit <return> to go on.')
commandwindow
pause

```

If the program had said `pause (2)`, the computer would have paused for 2 seconds before going on to the next programmed event, without waiting for input from the user. Non-integer values for `pause`, such as `pause (2.5)`, are possible, but beware that actual pause durations, whether they are triggered by integer or decimal values in the `pause` command, are imprecise owing to the nature of the `pause` command itself, not because of any inherent problem with MATLAB or, presumably, your computer.

6.4 Recording Reaction Times and other Delays With `tic ... toc`

Behavioral scientists often measure reaction time, the time for a response after some stimulus. Reaction time provides an index of decision-making. The longer the reaction time, the longer the component processes that led to it, all else being equal.

MATLAB provides a way of recording reaction times. The commands, which we introduced in Section 5.5, are called, appropriately, `tic` and `toc`. The `tic` command causes MATLAB to note the time when the `tic` command is issued. The `toc` command causes MATLAB to note the time that elapsed since the last `tic`. It is possible to measure reaction time by having people interact with the computer between `tic` and `toc`, as in this example.

Code 6.4.1:

```
commandwindow
tic
response = input('What is five plus the square root of 64? ')
Reaction_Time = toc
```

Output 6.4.1:

```
response =
    13
Reaction_Time =
    3.4589
```

The value returned by `toc`—in this case, the value of the variable called `Reaction_Time`—is expressed in seconds. Note that the synchronization of your display with the program and the speed of your computer’s keyboard can affect the precision of the value returned by `toc` in ways that can be difficult to assess, with uncertainties of up to several tens of ms (Plant & Quinlan, 2013; Plant & Turner, 2009). Other factors that may affect the accuracy of recorded times may be the model of your keyboard, mouse, or display. More precise timing is possible with a special application called Psychtoolbox, which is covered in Chapter 13. However, `tic . . . toc` may be sufficient if you are interested in long reaction times (half a second or so) that are large relative to the variability of the keyboard’s timing, provided you average across a number of trials, making the standard error of the mean sufficiently small (Ulrich & Giray, 1989). Running all your conditions with the same hardware can also make the timing data more comparable over conditions than they would be otherwise.

6.5 Formatting Numbers for Screen Outputs

When data are printed into the Command window, you can achieve some control of the form of the numerical output by using the `format` command. By typing `help format` or `doc format` in the MATLAB command line, you can learn about the options associated with the `format` command. Here are some of them.

Code 6.5.1:

```
t = [-.5:.5:1]';
format bank
bank_format_t = t
```

```

format compact
compact_t = t

format rat
rational_format_t = t

format short
short_format_t = t

format short g
short_g_format_t = t

format long
long_format_t = t

format long g
long_g_format_t = t

format loose
loose_t = t

format % return format to standard default
standard_format_t = t

```

Output 6.5.1:

```

bank_format_t =
    -0.50
     0
    0.50
    1.00

compact_t =
    -0.50
     0
    0.50
    1.00

rational_format_t =
    -1/2
     0
    1/2
     1

short_format_t =
    -0.5000
     0
    0.5000
    1.0000

```

```
short_g_format_t =
    -0.5
     0
    0.5
     1
long_format_t =
   -0.5000000000000000
           0
    0.5000000000000000
    1.0000000000000000
long_g_format_t =
           -0.5
           0
           0.5
           1

loose_t =
           -0.5
           0
           0.5
           1

standard_format_t =
   -0.5000
           0
    0.5000
    1.0000
```

6.6 Assigning Arrays of Literal Characters (Strings) to Variables

In the earlier examples of asking a user for his or her favorite number using `input` (Code 6.2.1), the user's response was interpreted as a number. MATLAB can be prompted to accept strings as input instead of numbers. In the code that follows, we indicate that a string should be accepted as input. To achieve this, we add a comma and 's' after 'What is your name? '. The 's' argument to the `input` function informs MATLAB that it should accept a string.

Code 6.6.1:

```
name = input('What is your name? ', 's')
```

Output 6.6.1:

```
What is your name? David
name =
David
```

It would be nice to reply to the user by name, but how can you do this without knowing what the user's name will be? `sprintf` is useful for this purpose. `sprintf`—short for

string print format—lets you assign data to a string variable. This sort of assignment is illustrated below, where we print Hello along with the string variable that follows. The percent sign tells MATLAB that the character following it is not part of the string to be printed, but rather denotes the format in which to print the variable as well as where to insert the variable into the string. The variable itself is indicated afterward.

Code 6.6.2:

```
name = input('What is your name? ', 's');
greeting = ...
    sprintf('Hello, %s, I will try to help you.', name);
greeting
```

Output 6.6.2:

```
What is your name? David
greeting =
Hello, David, I will try to help you.
```

In addition to %s, other formatting specifications can be used with sprintf:

- %d indicates that the next variable to be output will be an integer.
- %e indicates that the next variable to be output will be in scientific notation (e.g., 6.5e6, which is equal to 6.5×10^6 or 6.5 million).
- %f indicates that the next variable to be output will be a floating point (or decimal) number.
- \n indicates that a return will be included in a string.
- \t indicates that a tab character will be included in a string.

Examples follow.

Code 6.6.3:

```
piVal = sprintf('The approximate value of %s is %f', 'pi', pi)
```

Output 6.6.3:

```
piVal =
The approximate value of pi is 3.141593
```

Code 6.6.4:

```
first = 3.00;
second = 5.25;
int_vs_float = sprintf(['Here are two numbers, an integer,...
    ' %d, and a float, %f.'], first, second)
```

Output 6.6.4:

```
| int_vs_float =  
| Here are two numbers, an integer, 3, and a float,  
| 5.250000.
```

There are some other commands worth knowing about. `\n` in the format specification string indicates that a return will be included in the output of `sprintf`.

Code 6.6.5:

```
| twoLines = sprintf('two\nlines')
```

Output 6.6.5:

```
| twoLines =  
| two  
| lines
```

`%%` indicates that a percent sign (%) should be included in a string. Similarly, two apostrophes in a row, `' '`, indicate that an apostrophe should be included in a string. A single `%` or `'` would be misinterpreted as a format designator or the end of the formatting string, respectively.

Code 6.6.6:

```
| effort = sprintf(['Let's give %d%% effort' ...  
| ' to the project!!!'], 100)
```

Output 6.6.6:

```
| effort =  
| Let's give 100% effort to the project!!!
```

A final word about `sprintf` is that the presence of “print” within the word “sprintf” can be misleading. When you use the `sprintf` command, you are not actually printing in a physical sense. Rather, you are assigning data in string format (a sequence of literal, alphanumeric characters) to a variable.

A further indication that `sprintf` is not a command to physically print a variable is that in the examples above, each line of code that included the `sprintf` command lacked a semi-colon at the end of the line. The only property of the foregoing code that allowed the values to be displayed was that semi-colons were omitted from the ends of the lines. If you include a semi-colon at the end of a line that uses `sprintf` to assign its value to a string variable, MATLAB takes no observable action, though you can subsequently examine the value of the variable to check that the string variable was assigned a value, and that value can be re-used later.

Code 6.6.7:

```
| disp(effort)
```

Output 6.6.7:

```
| Let's give 100% effort to the project!!!
```

When a letter string is printed in the Command window by omission of the semi-colon, or by `disp`, it is left-justified, in contrast to when a number is printed, which is indented. The presence of indentation helps distinguish a number from a string composed of numeric characters, when they might be ambiguous.

Code 6.6.8:

```
| aNumber1234 = 1234
| aString1234 = '1234'
```

Output 6.6.8:

```
| aNumber1234 =
|             1234
| aString1234 =
| 1234
```

6.7 Controlling File Print Formats

We turn now to one of the most useful commands in MATLAB, `fprintf`, short for *file print format*. As its name suggests, `fprintf` lets you tailor the way your data are printed, just as `sprintf` lets you tailor the way your strings are constructed. We should tell you that if you expect only or mainly to shunt data to Excel spreadsheets—a very common need in behavioral science—you may not need to know the information that follows. Later, in Section 6.12, we will tell you how to put data into Excel spreadsheets and read them back.

Before showing examples of the `fprintf` command, it is useful to note that the special characters mentioned above in connection with `sprintf` also work with `fprintf`. For review purposes, those special characters are as follows:

- `%d` indicates that an upcoming variable will be printed as an integer.
- `%e` indicates that an upcoming variable will be printed in scientific notation (e.g., $6.5e6 = 6.5 \times 10^6 = 6.5$ million).
- `%f` indicates that an upcoming variable will be printed as a floating point (or decimal) number,
- `%s` indicates that an upcoming variable will be printed as a string.
- `\n` indicates that a return will be included in the printout.
- `\t` indicates that a tab character will be included in the printout.

Here are some examples that use `fprintf`. The first argument of `fprintf` is the format string, which indicates how the arguments that follow will be formatted. The format string can contain text that will appear in the output, but text that follows the `%` sign up until the next alphabetic character is special. It defines the format to be used for the output of the variables.

Code 6.7.1:

```
| fprintf('%s\n','Matlab can be fun.');
```

Output 6.7.1:

```
| Matlab can be fun.
```

Next, we print the matrix `[1:10]` first as integers (using `%d`), then in scientific notation (using `%e`), and finally in floating point notation (using `%f`). We print two returns after each line, one to end the line of print and one to add a line prior to the next one. If we don't tell MATLAB to print the returns, it won't do so.

Code 6.7.2:

```
| fprintf('%d',[1:10])  
| fprintf('\n\n')  
| fprintf('%e',[1:10])  
| fprintf('\n\n')  
| fprintf('%f',[1:10])  
| fprintf('\n\n')
```

Output 6.7.2:

```
| 12345678910  
  
| 1.000000e+002.000000e+003.000000e+004.000000e+005.000000e  
| +006.000000e+007.000000e+008.000000e+009.000000e+001.0000  
| 00e+01  
  
| 1.0000002.0000003.0000004.0000005.0000006.0000007.0000008  
| .0000009.00000010.000000
```

Output 6.7.2 isn't especially welcoming. It would be nice to have greater control of the output. The following example shows how to print the same matrix, `[1:10]`, specifying six characters for each value and treating each as a floating point number. The notation in quotes means, "Allocate six characters per number with two places to the right of the decimal point, using floating point notation."

Code 6.7.3:

```
fprintf('%6.2f', [1:10])
fprintf('\n')
```

Output 6.7.3:

```
1.00 2.00 3.00 4.00 5.00 6.00 7.00 8.00 9.00
10.00
```

The next example gives more information about how the formatting of numbers can be controlled. After defining `Pi_matrix` as a 1×10 matrix whose values are spaced linearly from `pi` to 2π , we print the values with no spaces to the right of the decimal point and then with two values to the right of the decimal point.

Code 6.7.4:

```
Pi_matrix = linspace(pi, 2*pi, 10);
fprintf('%6.0f', Pi_matrix);
fprintf('\n');
fprintf('%6.2f', Pi_matrix);
fprintf('\n');
```

Output 6.7.4:

```
3 3 4 4 5 5 5 6 6 6
3.14 3.49 3.84 4.19 4.54 4.89 5.24 5.59 5.93 6.28
```

The actual values of `Pi_matrix` are unaffected by the way they are printed. Even though the printed output is rounded to the nearest printable value, the original is unchanged, as can be confirmed by printing a sample of the original variable with different resolution.

Code 6.7.5:

```
format long
Pi_matrix(1:4)
```

Output 6.7.5:

```
ans =
3.141592653589793 3.490658503988659
3.839724354387525 4.188790204786391
```

Just as we can control the formatting of real numbers, it's also possible to control the format of integers. Here we allocate four characters per integer, then five characters per integer, and finally six characters per integer.

Code 6.7.6:

```
fprintf('%4d',[1:10]);  
fprintf('\n\n');  
fprintf('%5d',[1:10]);  
fprintf('\n\n');  
fprintf('%6d',[1:10]);
```

Output 6.7.6:

```
1  2  3  4  5  6  7  8  9 10  
  
1  2  3  4  5  6  7  8  9 10  
  
1  2  3  4  5  6  7  8  9 10
```

Note that we can insert tab characters between the printed values using the '`\t`' formatting operator. This is useful for generating output that may eventually be pasted into Excel or SPSS:

Code 6.7.7:

```
fprintf('%d\t',[1:10].^4);
```

Output 6.7.7:

```
1  16  81  256  625  1296  2401  4096  6561  10000
```

Even if a line containing `fprintf` ends with a semi-colon, printing still occurs. By contrast, as seen in the last section, `sprintf` assigns a string to a *variable*. For this reason, when an `sprintf` command is issued, the value of the variable is only printed if no semi-colon ends the line.

If different formats need to appear in a single line, those formats must be specified individually, so each value can be output as desired. Here is an example in which different formats are generated for different parts of each line of output.

Code 6.7.8:

```
a = [3.1:5.1];  
b = [3:5];  
c = a*2;  
d = b + 2;  
fprintf('%6.2f',a);  
fprintf('%4d',b);  
fprintf('\n');  
fprintf('%6.2f',c);  
fprintf('%4d',d);  
fprintf('\n');
```

Output 6.7.8:

3.10	4.10	5.10	3	4	5
6.20	8.20	10.20	5	6	7

It doesn't matter that the `fprintf` commands occupy different lines themselves. They could be arranged from left to right on one line separated by semi-colons and the output would be the same. (This is first time in this book we have mentioned the fact that different commands can be issued on the same line. Only for values in a matrix do line returns actually carry any meaning for MATLAB. We prefer line-by-line commands for visual clarity in most circumstances, but here, just to show what multi-command lines can look like, we show them.)

Code 6.7.9:

```
a = [3.1:5.1];
b = [3:5];
c = a*2;
d = b + 2;
fprintf('%6.2f',a); fprintf('%4d',b); fprintf('\n');
fprintf('%6.2f',c); fprintf('%4d',d); fprintf('\n');
```

Output 6.7.9:

3.10	4.10	5.10	3	4	5
6.20	8.20	10.20	5	6	7

Text and variables can be printed at the same time. Here we allocate two spaces (`%2d`) for the second variable printed, so the last column is right-justified even though some numbers have one digit and some have two.

Code 6.7.10:

```
a= [1 2 3 4 5];
asq = a.*a;
for i = 1:5
    fprintf('The square of %d is %2d\n',a(i),asq(i))
end
```

Output 6.7.10:

The square of 1 is	1
The square of 2 is	4
The square of 3 is	9
The square of 4 is	16
The square of 5 is	25

6.8 Writing Data to Named Files

`fprintf` can write data to files. In Code 6.8.1, which we will take a while to lead up to, we define a file into which data will be written using `fprintf`. We first open or create the file using the `fopen` command. The particular file in this example is named `mydata.txt`. Being a text file, the file has the suffix `.txt`. To enable writing to the file, we use `'wt'` as the second argument. The variable that is the output of `fopen`, which we called `fid`, is a file identifier. That value, `fid`, will be used in subsequent `fprintf` commands to direct the output to the file `mydata.txt`. There is nothing special about the name `fid`. We could just as well have called it `ham_and_eggs`.

Once we have assigned the file identifier to a file with `fid`, we can write data to it. There is no harm in also writing the data to the Command window using a separate `fprintf` command to make sure it looks the way we expect. In the program that follows, we write data to `fid` as well as the Command window, then we write more data to both locations, and finally we close `fid` using the `fclose` command. Until the file has been closed, it is not accessible for reading.

In the code that follows, besides doing the things already mentioned, we define a matrix called `rr` and print `rr` both to `fid` and the Command window. Note that we issue one print command at a time, first for `fid` and then for the Command window. The Command window is specified implicitly by omitting an output file name after the opening parenthesis following `fprintf`. The program ends by using the `type` command to report the contents of the file `mydata.txt` to verify that everything worked as intended.

Code 6.8.1:

```
fid = fopen('mydata.txt','wt');
rr = [1.1:5.1];

fprintf(['Data echoed to Command window as it is written'...
        ' to mydata.txt\n'])
fprintf('%6.1f',rr);           % to Command window
fprintf(fid,'%6.1f',rr);       % to file associated with fid

fprintf(fid,'\n');
fprintf('\n');

fprintf('%6.1f',rr+2);
fprintf(fid,'%6.1f',rr+2);
fprintf('\n\n')
fclose(fid);

fprintf('Data as read from mydata.txt:\n')
type mydata.txt
```

Output 6.8.1:

```
Data echoed to Command window as it is written to mydata.txt
    1.1    2.1    3.1    4.1    5.1
    3.1    4.1    5.1    6.1    7.1
```

```
| Data as read from mydata.txt:
|   1.1   2.1   3.1   4.1   5.1
|   3.1   4.1   5.1   6.1   7.1
```

There are other ways to write data to named files besides `fprintf`. One is to use `dlmwrite`. Here is an example in which the matrix data is saved as tab-delimited text to a file called `my_dlm_data`.

Code 6.8.2:

```
| data = [78:90];
| dlmwrite('my_dlm_data.txt',data,'\t');
| type my_dlm_data.txt;
```

Output 6.8.2:

```
| 78   79   80   81   82   83   84   85   86   87   88   89   90
```

For more information about `dlmwrite` and for pointers to other methods for writing data to named files, type `help dlmwrite` in the MATLAB Command window.

6.9 Writing Text to Named Files

The last section showed you how to write *data* to a file with the `fprintf` command. This section shows you how to write *text* to a file with `fprintf`. Here's an example.

Code 6.9.1:

```
| a = [1 2 3 4 5];
| acube = a.^3;
| myoutfile = fopen('CubesList.txt','wt');
| for i = 1:5
|     fprintf(myoutfile,'The cube of %d is %3d\n',a(i), ...
|         acube(i));
| end
| fclose(myoutfile);
| type('CubesList.txt');
```

Output 6.9.1:

```
| The cube of 1 is    1
| The cube of 2 is    8
| The cube of 3 is   27
| The cube of 4 is   64
| The cube of 5 is  125
```

6.10 Checking and Changing the Current Directory

The output listed above appeared in the Command window. Then we used `type` to confirm that the files `mydata.txt` and `my_dlm_data.txt` were saved as hoped. There are other ways to check for the existence of files. One is to list the contents of the current directory, using the `ls` command. Here is that command and the result obtained on the particular computer used to write this chapter. You should not expect to see all these files on your computer.

Code 6.10.1:

```
| ls
```

Output 6.10.1:

```
| CubesList.txt  my_dlm_data.txt  mydata.txt
```

The recently-created files, `mydata.txt` and `my_dlm_data.txt` are there.

Another way to list the current directory is with the `dir` command. In this example `dir` is used selectively, along with the `*` wildcard and the suffix that defines the file type (e.g., `.txt` or `.m`). Here is code that lists only the `.m` files in the current directory of the author at the time this example was written.

Code 6.10.2:

```
| dir *.m
```

Output 6.10.2:

```
| my_dlm_data.txt  mydata.txt
```

To find out the name of the current directory, you can use the `pwd` command. (Note that Mac OS uses forward slashes, `'/'`, instead of back slashes, `'\'`, to delimit folder names in directory listings and commands).

Code 6.10.3:

```
| pwd
```

Output 6.10.3:

```
| ans =  
| C:\Lab and Teach\PSU Teaching\Programming Seminar\Textbook
```

To change the current directory, you can use the `cd` command. To change to a specific named directory, its full path can be supplied, as in this example.

Code 6.10.4:

```
| cd('D:\MATLAB of David\')
| pwd
```

Output 6.10.4:

```
| ans =
| D:\MATLAB of David
```

To access the parent directory of the current directory, you can write

Code 6.10.5:

```
| cd('.')          % [or:]    cd ..\
| pwd
```

Output 6.10.5:

```
| ans =
| D:\
```

Moving to the parent directory in this way lets you move to a different sub-directory, such as *Exercises*, which is in the same directory as *Textbook*.

Code 6.10.6:

```
| cd(['\Lab and Teach\PSU Teaching\Programming Seminar\' ...
|     'Textbook'])
| pwd
| cd('..\Exercises')
| pwd
```

Output 6.10.6:

```
| ans =
| C:\Lab and Teach\PSU Teaching\Programming Seminar\
| Textbook
| C:\Lab and Teach\PSU Teaching\Programming Seminar\
| Exercises
```

A command like `cd ..\Exercises` would also work, as long as there are no illegal characters such as spaces in the folder names specified. The current directory can also be changed by browsing, using your mouse, in the Current Folder window.

Changing the current directory can be useful for accessing data in different directories or for writing data to different directories.

6.11 Reading Data Saved as Plain Text From Named Files

How can data be read into a program from an external file? One way is to use the `load` command. You can use this command to load numerical data that have an equal number of entries per line. The name of the file to be loaded must be enclosed in single quote marks, within parentheses. It is easy to forget to include the single quote marks and then be frustrated by error messages that say no such file exists, so be careful about this.

Code 6.11.1:

```
| data_from_file = load('mydata.txt')
```

Output 6.11.1:

```
| data_from_file =  
|      1.1000    2.1000    3.1000    4.1000    5.1000  
|      3.1000    4.1000    5.1000    6.1000    7.1000
```

You can also use `load` to read files in plain text format that may have been created with other programs, such as Microsoft Word. Be sure to save the files in plain text format if you plan to load them.

Another way to read files in plain text, if they are not purely numeric, is with the `fgetl` command; that last character is *el*, not *one*. This command reads files one line at a time into a matrix of characters (a string). If `fgetl` reads the last line of a file, the Boolean `feof` is set to `true`, as shown in Code 6.11.2, which is constructed to read until the end of the file is detected, echoing the lines to the Command window.

Code 6.11.2:

```
| myinfile = fopen('cubeslist.txt');  
| nlines = 0;  
| while true  
|     thisline = fgetl(myinfile);  
|     nlines = nlines + 1;  
|     fprintf('Line %d:  %s\n',nlines,thisline);  
|     if feof(myinfile)  
|         disp('all done!')  
|         break  
|     end  
| end
```

Output 6.11.2:

```
| Line 1:  The cube of 1 is    1  
| Line 2:  The cube of 2 is    8
```

```

Line 3:  The cube of 3 is  27
Line 4:  The cube of 4 is  64
Line 5:  The cube of 5 is 125
all done!

```

6.12 Reading Data From and Writing Data to Excel Spreadsheets

Reading data from Microsoft Excel files is easy in MATLAB, as is writing to such files. Here is how you can read an Excel spreadsheet called 'data' into a matrix, M.

Code 6.12.1:

```
| M = xlsread('data.xls');
```

You can also specify a particular worksheet to be read within the Excel document. In this case, the name of that worksheet, as previously saved in Excel, is 'Experiment 2'.

Code 6.12.2:

```
| M = xlsread('data.xls', 'Experiment 2');
```

`xlsread`, by default, will return only the numeric portion of the spreadsheet. Non-numeric cells, by default, will be assigned the value NaN (see Section 4.6), and column names will be ignored. You can go beyond this default mode, however, by taking advantage of the fact that `xlsread` can return three results, as indicated via `help xlsread`, namely `[num, txt, raw]=xlsread(FILE, RANGE)`. If you want just the text from the spreadsheet called 'Experiment 2' in `data.xls`, you can get it by printing out `TXT` (or whatever name you give to the second argument within the brackets to the left of the equal sign). If you want text and data, you can get it by printing out `RAW` (or whatever name you give to the third argument within the brackets to the left of the equal sign). If you want just the numbers from the spreadsheet called 'Experiment 2' in `data.xls`, you can get it by printing out `NUM` (or whatever name you give to the first argument within the brackets to the left of the equal sign), or you can name a single matrix, such as `M` in Code 6.12.2 above. By default, when just a single output is specified, it is the first of the set that can be gotten with more than one requested output.

Writing data to an Excel spreadsheet is easy. The relevant function is `xlswrite`. In Code 6.12.3 `xlswrite` is used to write `M` to an Excel file called `My_Excel_File`. (On Macintosh systems, MATLAB will write a comma-separated-value, or `.csv` file, rather than a `.xls` file).

Code 6.12.3:

```
| xlswrite('My_Excel_File', M);
```

You can specify a spreadsheet number as an optional value after the name of the file being exported to Excel.

6.13 Taking Precautions Against Overwriting Files

You may wish to know if the file you are about to write to already exists so you don't overwrite it. MATLAB's `exist` function can be used for this purpose. Code 6.13.1 shows how MATLAB can be used to test for the existence of a file called `filename` and give a warning if that filename is already taken. Note that the check is only made in the current directory. In this case, there are no quotes around `filename` in `if ~exist(filename)` because `filename` is a variable that contains the name as a string. It is not the name itself.

Code 6.13.1:

```
filename = input('File name: ', 's');
if ~exist(filename)
    dlmwrite('my_dlm_data.txt', data, '\t');
else
    disp(['Error: the file '' filename '' already exists!']);
end
```

Output 6.13.1:

```
File name: my_dlm_data.txt
Error: the file 'my_dlm_data.txt' already exists!
```

If a checked file has a suffix, it needs to be included in the test. For example, if `exist('mydata')` fails to yield a warning but you know the file is there, it may be that you need to say `exist('mydata.txt')`. Otherwise you may overwrite `mydata.txt`, having falsely concluded it is absent. These matters should be checked during program development, before you put your program to full use.

Another good strategy is to make sure every file you use has a unique filename. A way to do this is to make up a filename with a timestamp that reminds you of when the file was created. Here is an example using the `clock` command, which reads the full time of day into a 1×6 matrix, with the values representing the year, month, day, hour, minute, and seconds (including fractional seconds). Another way to make a data filename unique is to include the initials of a participant (say, Elmer Fudd) and the date (e.g., August 31, 2013).

Code 6.13.2:

```
timeofday = clock;
FirstOutputfilename = ...
sprintf('Expt5_%02d_%02d_%02d_T%02d%02d%02d.txt', ...
    round(timeofday(1:6)))
inits = input('Subject initials: ', 's');
SecondOutputfilename = strcat('Expt5_', inits, '_', ...
    sprintf('%02d_%02d_%02d', timeofday(1:3)), '.txt')
```

Output 6.13.2:

```

FirstOutputfilename =
Expt5_2013_08_25_T164855.txt
Subject initials: EF
SecondOutputfilename =
Expt5_EF_2013_08_25.txt

```

6.14 Saving and Loading Variables in Native MATLAB Format

Often, complex analysis is best conducted by running one program to organize the data for analysis and a second program to summarize the data across conditions. Approaching data analysis this way makes each individual step easier to design and evaluate. However, there needs to be an easy way to convey the output of one step of the analysis to the next, especially if the earlier step consumes considerable computational time or interaction by the researcher. The commands `save` and `load` let you take a snapshot of one or more variables in one step of a multistep process and pass those values on to the next step.

For the moment, assume that the data you are working with at the end of a first program step are in two variables, `rawdata` and `summarydata`. You need to know only the values of `summarydata` to run the program that executes the second analysis step, but you don't want to repeat the work of the first program while developing the second. Here, using `save` to generate a file of type `.mat` can be useful. Of all the ways of saving data from MATLAB computations, using `.mat` files has the advantage that the variables will be reloaded exactly as they were in the saving program, and the time required for writing and reading the data is the shortest. The limitation of `.mat` files is that they can be read only by MATLAB. They are meaningless to other programs.

Code 6.14.1:

```

% Step1Program.m
rawdata = load('mydata.txt');
% Perform the analysis to convert raw to summary data
% here
save('DatafromStep1.mat', 'summarydata')
whos

```

Output 6.14.1:

Name	Size	Bytes	Class
FirstOutputfilename	1x26	52	char
Firstoutputfilename	1x26	52	char
SecondOutputfilename	1x20	40	char
a	1x6	48	double
summarydata	23x33	6072	double
ans	1x20	40	char

b	1x24	48	char
d	2x5	80	double
filename	1x15	30	char
inits	1x2	4	char
mydata	2x5	80	double
outputfilename	1x20	40	char
rawdata	2x5	80	double
timeofday	1x6	48	double

Code 6.14.2:

```
% Step2Program.m
clear
load('DatafromStep1.mat')
whos
% Now perform the analysis on the summary data.
```

Output 6.14.2:

Name	Size	Bytes	Class
summarydata	23x33	6072	double

6.15 Learning More About Input and Output

MATLAB has more functions for input and output. For reading tabular data that includes text fields, `textscan` can extract columns using a pattern-matching syntax that is similar to that of `fprintf`. If the data are unstructured, it may be necessary to read each line individually with `fgetl` and do processing on a line-by-line basis. MATLAB can also deal with binary data, as obtained from scientific instruments, using `fread`, `fwrite`, and `fseek`. It is worth checking the documentation as well as the MathWorks website before writing new code that uses these commands. Typing `help iofun` can be informative in this regard. The `iofun` command provides a portal to all the material that has been covered here, plus more.

Code 6.15.1:

```
| Help iofun
```

Output 6.15.1:

```
File input/output.

File import/export functions.
    dlmread      - Read delimited text file.
    dlmwrite     - Write delimited text file.
    load         - Load workspace from MATLAB (.mat) file.
    save         - Save workspace or variables to MATLAB
                  (.mat) file
```

```

importdata - Load workspace variables disk file.
wklread    - Read spreadsheet (.wkl) file.
wklwrite   - Write spreadsheet (.wkl) file.
xlsread    - Read spreadsheet (.xls) file.

```

6.16 Practicing Input-Output

Try your hand at the following exercises, using only the methods introduced so far in this book or information given in the problems themselves.

Problem 6.16.1:

Write a program that yields the output shown below. Note that each element of B is the corresponding element of A, squared. Each value appears in the output with seven columns per number and with one place to the right of the decimal point. The output should look like this:

```

A
  1.0    2.0    3.0    4.0    5.0    6.0    7.0    8.0
9.0   10.0
 11.0   12.0   13.0   14.0   15.0   16.0   17.0   18.0
19.0   20.0
 21.0   22.0   23.0   24.0   25.0   26.0   27.0   28.0
29.0   30.0
 31.0   32.0   33.0   34.0   35.0   36.0   37.0   38.0
39.0   40.0
 41.0   42.0   43.0   44.0   45.0   46.0   47.0   48.0
49.0   50.0
 51.0   52.0   53.0   54.0   55.0   56.0   57.0   58.0
59.0   60.0

B
  1.0    4.0    9.0   16.0   25.0   36.0   49.0   64.0
81.0  100.0
 121.0  144.0  169.0  196.0  225.0  256.0  289.0  324.0
361.0  400.0
 441.0  484.0  529.0  576.0  625.0  676.0  729.0  784.0
841.0  900.0
 961.0 1024.0 1089.0 1156.0 1225.0 1296.0 1369.0 1444.0
1521.0 1600.0
 1681.0 1764.0 1849.0 1936.0 2025.0 2116.0 2209.0 2304.0
2401.0 2500.0
 2601.0 2704.0 2809.0 2916.0 3025.0 3136.0 3249.0 3364.0
3481.0 3600.0

```

Now make the output of B “tab-delimited,” so you could copy from the Command window and paste into an Excel or SPSS spreadsheet.

Problem 6.16.2:

Write a program that creates an Excel file that will serve as the spreadsheet into which data from a behavioral science experiment can be saved. The Excel file should have the following columns in each of 200 rows:

Column 1: subject_number (1 to 200)

Column 2: subject_number_parity (odd, denoted 1; or even, denoted 0)

Column 3: NaN, serving as a placeholder for the response to be given.

Column 4: NaN, serving as a placeholder for the accuracy of the response.

Column 5: A random value drawn from a normal distribution with mean equal to 0 and standard deviation equal to 1 for odd-numbered subjects, or a random value drawn from a normal distribution with mean equal to 10 and standard deviation equal to 5 for even-numbered subjects.

Read the Excel file back into MATLAB to observe the effects of having inserted NaN in columns 4 and 5.

Problem 6.16.3:

Write a program in which a user is asked for a password. The program should check whether the password is contained in a list of three acceptable six-letter passwords, each of which begins with a letter, defined as follows:

```
correct_passwords = ['A1B2C3'; 'B2C3A1'; 'C3A1B2']
```

Idiot-proof the program so the user is not rejected prematurely if he or she makes a typing error (e.g., too many or too few characters), but only let the user respond to the input a set number of times (e.g., 4).

Problem 6.16.4:

Modify the program from Problem 6.16.4 so passwords consist of six-digit numbers from 100,000 to 999,999, and the matrix of passwords is retrieved from an external file. You will need to create the external file first. Set it up so there are 100 passwords for 100 employees. Later, for an employee to enter the system, the password he or she supplies must be the password associated with his or her employee number, which is 1 through 100. Two pieces of information will help you solve this fairly difficult problem. One is that you can generate a 100×1 matrix of passwords from 100,000 to 999,999 as follows:

```
number_of_employees = 100;  
passwords = randi(899999,number_of_employees,1)+100000
```

Second, you will need to convert the password number entered by the user to a number from a string. You can achieve this conversion with a command that will be officially premiered in Chapter 7, `str2num`. The following code segments will also be useful. Note that the `if` statement need not immediately follow the `input` statement in your program. If you omit the 's', the input will already be a number.

```

yourpassword = input('What is your 6 digit password? ', 's')
if passwords(employee_number,:) == str2num(yourpassword)
    disp([OK_to_enter])
end

```

Problem 6.16.5:

One format of data files that can be imported into Excel or SPSS has the following characteristics: The first line of such a file is a header file, a series of valid SPSS variable names, tab delimited. Subsequent lines are numerical with the subject number in the first column followed by the scores for that subject in subsequent columns (i.e., also tab delimited). Generate output to the Command window that describes data for six subjects and two conditions (call the conditions 'left' and 'right'). The header line will then read,

```
subno left right
```

and the first data line (second line printed out) will be

```
1      0.32  0.54
```

if 0.32 and 0.54 are the two scores for subject 1. Print such a data set in the Command window using `fprintf` and verify that you can copy and paste the data set into Excel. Then modify your code to write into a data file (`handednessdata.txt`) with the same contents, rather than the Command window. Verify that you can open the file with Excel and that all the numbers end up in the right places.

You can generate your dataset by `thedata = [rand(6,2)]`. The resulting matrix will have six rows (one for each subject) and two columns (the left and right score for each participant).

Problem 6.16.6:

Use the `fprintf` command to write a limerick or *haiku*, on the topic of MATLAB programming, appropriately formatted, to the file `mypoetry.txt`. Verify the content by `type mypoetry.txt`.

Problem 6.16.7:

Make an array using `magic(N)`, where N can be any value between 3 and 9. Print out the array, along with the row and column sums (the marginal sums) in a table formatted like the one below. Write the program in a sufficiently general fashion that it would work for any square array, not just the one you tested. Test your program on `M = randi(9,N,N)` as well as on `M = magic(N)`;

Marginal sums for $N = 5$

```

17 24  1  8 15 | 65
23  5  7 14 16 | 65
 4  6 13 20 22 | 65
10 12 19 21  3 | 65
11 18 25  2  9 | 65
-- -- -- -- --
65 65 65 65 65

```

7. Data Types

This chapter covers the following topics:

- 7.1 Identifying strings, numbers, and logical values (Booleans)
- 7.2 Converting characters to numbers and vice versa
- 7.3 Creating, accessing, and using cell arrays
- 7.4 Creating and accessing structures
- 7.5 Searching and modifying strings
- 7.6 Applying data types
- 7.7 Practicing data types

The commands that are introduced and the sections in which they are premiered are:

<code>class</code>	(7.1)
<code>double</code>	(7.1)
<code>single</code>	(7.1)
<code>char</code>	(7.2)
<code>num2str</code>	(7.2)
<code>str2num</code>	(7.2)
<code>strcat</code>	(7.2)
<code>{ }</code>	(7.3)
<code>cell2mat</code>	(7.3)
<code>deal</code>	(7.4)
<code>record.field</code>	(7.4)
<code>feof</code>	(7.5)
<code>fgetl</code>	(7.5)
<code>strcmp</code>	(7.5)
<code>strcmpi</code>	(7.5)
<code>strfind</code>	(7.5)
<code>strrep</code>	(7.5)
<code>textscan</code>	(7.5)

7.1 Identifying Strings, Numbers, and Logical Values (Booleans)

In previous chapters you were exposed to different types of data: numbers (including matrices of numbers) and strings (sequences of alphabetic or numeric characters not directly usable in numerical calculations). These are just two of the types of data representation used in MATLAB. A fuller list is provided below via code designed to spawn many, though not all, of the data types to be introduced in this chapter. One thing to note about the data

types, whose identities are revealed through the `whos` command, is that numbers can be of different types.

Code 7.1.1:

```
clear all

a = 'a'
b = 1
c = 1.00
d = round(c)
e = single(d)
f = uint8(e)
g = true
h = false
i = 'abcde'
j = 'ABCDE'
k = i - j
l = int8(k)

whos
```

Output 7.1.1:

```
a =
a
b =
    1
c =
    1
d =
    1
e =
    1
f =
    1
g =
    1
h =
    0
i =
abcde
j =
ABCDE
k =
    32    32    32    32    32
l =
    32    32    32    32    32
```


Name	Size	Bytes	Class	Attributes
a	1x1	2	char	
b	1x1	8	double	
c	1x1	8	double	
d	1x1	8	double	
e	1x1	4	single	
f	1x1	1	uint8	
g	1x1	1	logical	
h	1x1	1	logical	
i	1x5	10	char	
j	1x5	10	char	
k	1x5	40	double	
l	1x5	5	int8	

As seen above,

a, which was set to 'a', is an array of type `char` and uses two bytes of memory;

b, which was set to 1, is an array of type `double` and uses eight bytes of memory;

c, which was set to 1.00, is an array of type `double` and also uses eight bytes of memory;

d, which was set to `round(c)`, is an array of type `double` and again uses eight bytes of memory;

e, which was set to `single(c)`, is an array of type `single` and uses just four bytes of memory;

f, which was set to `uint8(e)`, is an array of type `uint8` (an unsigned integer, eight bits long) and uses just one byte of memory;

g and h, which were set to logical or “Boolean” values (after the British logician George Boole) are each arrays of type `logical` (i.e., true or false) and use just one byte of memory;

i and j, which were set to five-character strings of type `char`, each use 10 bytes of memory;

k, which shows the difference between i and j, is an array of type `double`, and uses 40 bytes of memory (eight bytes for each `double` value);

l, the letter “el,” not the number 1, which shows the same value as k, but in an array of type `int8`, and uses five bytes of memory. The integer value of 32 fits within the range limits of an 8-bit signed integer (–128 to 127).

These examples reveal several features of the data types represented. One is that in MATLAB, a number is, by default, a `double`, a value stored with “double value precision,” requiring 8 bytes of memory, taking on a value between -2×10^{308} and $+2 \times 10^{308}$. Another feature is that even when a `double` is rounded, it takes 8 bytes of memory. This is true even if the number is passed through `floor` or `ceil` (see Chapter 4). A third feature is that

a number can be a `single` (i.e., a number stored with single value precision, requiring just four bytes of memory), or one of several integer types that require just one byte of memory. A fourth is that variables can be assigned to the type `logical`, whose possible values are 1 and 0. In this context, 1 means the same as `true` and 0 means the same as `false`.

Why is it helpful to know about data types? One reason is that different data types require different amounts of memory. A variable of type `double` requires more memory than a variable of type `single`. This can be important if your program is memory-intensive, as may be the case if it uses a great many variables or very large data sets.

It is easy to convert values of one data type to another. The possible conversion commands can be found with `help datatypes`. The output that results is more complete than what follows, but the material below is likely to be instructive.

Code 7.1.2:

```
| help datatypes
```

Output 7.1.2:

```
Data types and structures.

Data types (classes)
double             - Convert to double precision.
logical            - Convert numeric values to logical.
cell               - Create cell array.
struct             - Create or convert to structure array.
single             - Convert to single precision.
uint8              - Convert to unsigned 8-bit integer.
uint16             - Convert to unsigned 16-bit integer.
uint32             - Convert to unsigned 32-bit integer.
uint64             - Convert to unsigned 64-bit integer.
int8               - Convert to signed 8-bit integer.
int16              - Convert to signed 16-bit integer.
int32              - Convert to signed 32-bit integer.
int64              - Convert to signed 64-bit integer.
```

You can learn the data type of a variable by using the `class` function.

Code 7.1.3:

```
double_value = 2
class(double_value)

single_value = single(double_value)
class(single_value)
```

Output 7.1.3:

```
| double_value =
      2
```

```
ans =  
double  
  
single_value =  
    2  
ans =  
single
```

Another reason to know about data types is that the range of possible values differs for different types of values. For example, the 256 possible values of a variable of type `int8` range from -128 to 127 , whereas the 256 values of a `uint8` range from 0 to 255 . By contrast, the maximum precision that can be represented in a variable of type `double` is 15 significant digits. Larger values can be represented using scientific notation, but some precision may be lost due to rounding of values with more than 15 significant digits, such as values greater than 9×10^{15} . Consider the following example:

Code 7.1.4:

```
sum = 0  
while sum ~= 1  
    fprintf('Not there yet...\n')  
    sum = sum + 1/99;  
    fprintf('%f\n', sum)  
end
```

The output is omitted because it is infinitely long. A small variation will show why:

Code 7.1.5:

```
sum = 0  
while sum ~= 1  
    fprintf('Not there yet...\n')  
    sum = sum + 1/99;  
    fprintf('%17.15f\n', sum)  
    if sum > 1.05  
        break  
    end  
end
```

Output 7.1.5:

```
[... 95 lines of output omitted]  
Not there yet...  
0.969696969696968  
Not there yet...  
0.979797979797978  
Not there yet...  
0.989898989898988  
Not there yet...
```

```

0.9999999999999998
Not there yet...
1.010101010101008
Not there yet...
1.020202020202018
Not there yet...
1.030303030303028
Not there yet...
1.040404040404039
Not there yet...
1.050505050505049
[And so on, ad infinitum...]

```

As you have just seen, because $1/99$ is an infinitely repeating decimal (0.0101010101010...) the sum of 99 terms is not exactly 1.0. The value of `sum` that falls closest to 1 in the computer's notation is literally one *bit* too small (0.999999999999998) due to rounding error. See Hayes (2012) for further discussion of precision issues in numerical computing.

A final reason to know about data types is that this knowledge can help you gain greater control over the speed with which your computer can communicate with external equipment in experiments you may conduct. Such communication typically requires the use of MATLAB's Data Acquisition Toolbox. More will be said about toolboxes later in this book. Typically, such communication uses `int8` or `uint8` variables rather than `double` variables because it takes less time to transmit 1 byte of information than the 8 required by a `double`.

7.2 Converting Characters to Numbers and Vice Versa

In computers, all data are ultimately represented as binary digits (or “bits,” for short) that make up numbers, so alphabetic characters and other symbols, such as exclamation marks, can be expressed in terms of their numerical equivalents. The code below shows how to get the numerical equivalents of characters using the `double` function.

Code 7.2.1:

```

de = double('!!')
dq = double('Let's go!')

```

Output 7.2.1:

```

de =
    33
dq =
    76    101    116    39    115    32    103    111    33

```

As seen above, `double` gives the matrix of numbers associated with the string of characters. To reverse the operation, `char` gives the characters associated with numbers. The program below show the character equivalents of the numerical matrices `de` and `dq`.

Code 7.2.2:

```
de_lettered = char(de)
dq_lettered = char(dq)
```

Output 7.2.2:

```
de_lettered =
!
dq_lettered =
Let's go!
```

Converting between characters and numbers can be useful in behavioral science. An example application would be expressing categorical responses as numbers if your data analysis justifies that procedure. Your data would be more readable if you code participants' sex as 'm' or 'f' than as 0 (for male) and 1 (for female), say.

There is a consistent relation between character codes and typed characters, based on the ASCII (American Standard Code for Information Interchange) standard. For the mono-spaced Courier font used by MATLAB, a subset of the code equivalents can be generated as follows:

Code 7.2.3:

```
for i = 1:8
    for j = (i:11:91)
        thiscode = j+31;
        fprintf('%5.0f %s',thiscode,char(thiscode));
    end
    fprintf('\n');
end
```

Output 7.2.3:

```
32      43 +   54 6   65 A   76 L   87 W    98 b   109 m   120 x
33 !    44 ,    55 7   66 B   77 M   88 X    99 c   110 n   121 y
34 "    45 -    56 8   67 C   78 N   89 Y   100 d   111 o   122 z
35 #    46 .    57 9   68 D   79 O   90 Z   101 e   112 p
36 $    47 /    58 :    69 E   80 P   91 [   102 f   113 q
37 %    48 0    59 ;    70 F   81 Q   92 \   103 g   114 r
38 &    49 1    60 <    71 G   82 R   93 ]   104 h   115 s
39 '    50 2    61 =    72 H   83 S   94 ^   105 i   116 t
```

Another thing to keep in mind is that the human readable version of a number is of type `string`. By contrast, the computer-readable version is of one of the numerical types. You can easily convert from one to the other using the commands `num2str` and `str2num`, as in this example.

Code 7.2.4:

```
num1 = 123.456
str1 = '567.890'
```

```
strOfNum1 = num2str(num1)
numOfStr1 = str2num(str1)
whos
```

Output 7.2.4:

```
num1 =
    123.4560
str1 =
    567.890
strOfNum1 =
    123.456
numOfStr1 =
    567.8900
```

Name	Size	Bytes	Class	Attributes
num1	1x1	8	double	
numOfStr1	1x1	8	double	
str1	1x7	14	char	
strOfNum1	1x7	14	char	

Sometimes you need to both convert numbers to strings, and join (concatenate) the strings. Here is an example of a situation where these two needs arise. Beware that the example will first be presented to you in the form of code that yields an error message.

Code 7.2.5:

```
fave = 7;
disp('Your favorite number is ' fave);
```

Output 7.2.5:

```
??? Error: File: Number_To_String_01.m Line: 4 Column: 31
Missing MATLAB operator.
```

Why did MATLAB return an error message? The reason is that `disp` requires a single matrix, consisting of one row of values that need to be of one type, either all numbers or all alphabetic characters. The strings to be printed can be concatenated using brackets or using the `strcat` command. Note, however, that `strcat` ignores trailing spaces in any of the concatenated strings.

Code 7.2.6:

```
fave = 7;
disp(['Your favorite number is ' int2str(fave) '.']);
disp(strcat('Your favorite number is ', int2str(fave), '.'));
```

Output 7.2.6:

```
| Your favorite number is 7.  
| Your favorite number is7.
```

A particularly useful variant of `num2str` can be used to print numbers in tab-delimited form, by including a formatting string in the command. Output generated this way is easy to copy from the Command window to a spreadsheet.

Code 7.2.7:

```
| disp(num2str([1:5].^2, '%d\t'))  
| disp(num2str([1:5].^.5, '%g\t'))
```

Output 7.2.7:

```
| 1      4      9     16     25  
| 1      1.41421    1.73205      2      2.23607
```

7.3 Creating, Accessing, and Using Cell Arrays

If you look back at Output 7.1.2, you will see mention of a data type that has not been referred to before in this book. That data type is the `cell`. A cell is an array with the convenient property that each of its elements can store a matrix of a different size or type—a single number, a numerical matrix, or a string. The reason this is a convenient property is that you may sometimes need to represent variables of different sizes. For example, if you are doing a study with a list of words, where the number of letters differs for the words, as in 'apples' (6 letters) or 'oranges' (7 letters), you can't put the apples and oranges into rows of the same matrix. In the following code, we do so, however, just to show that, by not yet incorporating `cell`, you “upset the applecart.” Here, `MyMatrix1` gets an array of two strings of the same length, and `MyMatrix2` gets a mixed array in which the rows have different numbers of letters. It should come as no surprise that MATLAB balks at the assignment when the row lengths differ.

Code 7.3.1:

```
| MyMatrix1 = [  
| 'oranges'  
| 'bananas']  
  
| MyMatrix2 = [  
| 'apples'  
| 'oranges']
```

Output 7.3.1:

```
| MyMatrix1 =  
| oranges  
| bananas
```

```
Error using vertcat
Dimensions of matrices being concatenated are not consistent.
```

It can be frustrating to think that if you are creating an array of pigeon-holes, so to speak, it's only possible to have the same size of hole in every row and column. That constraint is “for the birds!” Cell arrays let you circumvent this problem.

To use cell arrays instead of matrices, use braces (“curly brackets”) rather than square brackets, as in the code below.

Code 7.3.2:

```
MyCells = {
    'apples'
    'oranges'}
for i = 1:2
    thisword = MyCells{i}
end
whos
```

Notice that MATLAB indicates that `MyCells` is a 2×1 cell array; `apples` is in cell 1 of the array, and `oranges` is in cell 2. You can address the rows and column of a cell array using braces, having them serve the same function as do parentheses in addressing matrices. You can address the individual elements of the cell array using `Mywords{i}`, which returns a string in this case, because each element in `Mywords` is a string.

Output 7.3.2:

```
MyCells =
    'apples'
    'oranges'
thisword =
apples
thisword =
oranges
```

Name	Size	Bytes	Class	Attributes
MyCells	2x1	250	cell	
i	1x1	8	double	
thisword	1x7	14	char	

The semi-colon within braces concatenates rows vertically in a cell array just as it does within brackets for the rows of a numerical matrix. The following code also shows that the elements of a matrix (string or numeric) in a cell array can be further addressed by putting the desired index in parentheses after the index in braces that selects the particular cell.

Code 7.3.3:

```
c = {[ 1 2 3]
     [4 5 6 7]
     ['rats mice']; [' voles']
     [1 3]}
c_second_row = c{2}
c_second_row_middle_numbers = c{2}(2:3)
c_third_row = c{3}
c_third_row_second_character = c{3}(2)
```

Output 7.3.3:

```
c =
     [1x3 double]
     [1x4 double]
     'rats mice'
     ' voles'
     [1x2 double]
c_second_row =
     4     5     6     7
c_second_row_middle_numbers =
     5     6
c_third_row =
rats mice
c_third_row_second_character =
a
```

As seen above, cells are not only useful for representing arrays with different number of rows or columns of a given data type; they are also useful for representing arrays of different data types, such as strings and numbers. This point was illustrated above without specifically mentioning it, but if you were paying close attention, you might have exclaimed while looking at Code 7.3.3, “Wow, the elements of cell array *c* include both numbers *and* words!”

Here is another example of a cell array, called `Names_and_Numbers`, whose entries have different lengths and are of different types. To access individual values within `Names_and_Numbers`, you can use `cell2mat`. Note that in Code 7.3.4, the opening brace must appear on the same line as `=`. As seen in Output 7.3.4, `cell2mat` not only converts numbers within cells to doubles; it also converts strings within cells to character strings.

Code 7.3.4:

```
Names_and_Numbers = {
'Bob' [90 95]
'Jane' 100
}
```

```
Name1 = cell2mat(Names_and_Numbers(1,1))
Numbers1 = cell2mat(Names_and_Numbers(1,2))
```

Output 7.3.4:

```
Names_and_Numbers =
    'Bob'      [1x2 double]
    'Jane'     [      100]
Name1 =
Bob
Numbers1 =
    90     95
```

Having obtained the contents of the cell, you can make use of it as you do with other kinds of data. Cell arrays can be used as the control variables in `for` and `switch` statements, similar to numerical matrices. In this `for` loop, the operations are repeated once for each cell in the array. The contents of the cell have to be converted to a character string or matrix for further computation in the loop.

Code 7.3.5:

```
for produce = {'Apple' 'Artichoke' 'Banana' 'Broccoli'...
              'Cherry' 'Cauliflower'}
    productName = char(produce); % convert cell to char
    switch productName
        case {'Apple' 'Banana' 'Cherry'}
            fprintf('%s is a fruit.\n', productName);
        case {'Artichoke' 'Broccoli' 'Cauliflower'}
            fprintf('%s is a vegetable.\n', productName);
    end
end
```

Output 7.3.5:

```
Apple is a fruit.
Artichoke is a vegetable.
Banana is a fruit.
Broccoli is a vegetable.
Cherry is a fruit.
Cauliflower is a vegetable.
```

The following example, which computes the number of days in a month, is based on suggestions by Henk Heijink, who was a graduate student when the first edition of this book was prepared, and also by Christopher Stevens, who was a graduate student when the second edition of this book was prepared. It uses a cell array to represent the months in one of the `case` statements.

Code 7.3.6:

```
% Days_In_A_Month
month = input('Type in the month: ','s');
year = input('Type in the year (4 digits): ');
switch month
    % Thirty days hath September,
    % April, June, and November...
    case {'September' 'April' 'June' 'November'}
        no_of_days = 30;
    case 'February'
        if rem(year, 4) == 0 & ...
            (rem(year, 100) ~= 0 | rem(year, 400) == 0)
            no_of_days = 29;
        else
            no_of_days = 28;
        end
    % All the rest have thirty-one
    otherwise
        no_of_days = 31;
end
fprintf('%s %d has %d days.\n', month, year, no_of_days);
```

Output 7.3.6:

```
Type in the month: February
Type in the year (4 digits): 2012
February 2012 has 29 days.
```

Note that each case may be selected by more than one value of the switch variable (case {'September' 'April' 'June' 'November'}), and that a particular case can have many lines of code in its implementation, as does the case for 'February'.

7.4 Creating and Accessing Structures

In Code 7.3.4, the two columns of the cell array represented different aspects of the data: subject names in column 1 and a numeric matrix in column 2. When you are dealing with this kind of data representation, you need to remember that column 1 has the *names* and column 2 has the *numbers*. This is not too hard to remember if there are only two columns, but it could get hard if you had many variables to keep track of or if you were sharing the code with a colleague. It would be useful to have a representation that facilitates the identification of variables so you don't have to keep the identification rule in mind or explicitly comment it. Here is how you can represent the data of Code 7.3.4 using a special data type in MATLAB, the structure, or `struct`.

Code 7.4.1:

```
Names_and_Numbers(1).name = 'Bob';
Names_and_Numbers(2).name = 'Jane';
Names_and_Numbers(1).RTs = [90 95];
Names_and_Numbers(2).RTs = [100];
Names_and_Numbers
```

Output 7.4.1:

```
Names_and_Numbers =
1x2 struct array with fields:
    name
    RTs
```

The struct variable `Names_and_Numbers` has two elements (1 and 2), each of which has two fields (`name` and `RTs`). You can apply this way of coding information in a behavioral experiment in which three stimulus factors (`side`, `intensity`, and `duration`) vary from trial to trial. The first trial presents a left, bright stimulus lasting 200 ms; the second trial presents a right, bright stimulus lasting 300 ms; and the third trial presents a left, dim stimulus lasting 400 ms. One approach, which does not use a struct, is to define a matrix called `trials` whose first column specifies side (1 = left, 2 = right), whose second column specifies brightness (1 = dim, 2 = bright), and whose third column specifies stimulus duration.

Code 7.4.2:

```
trials = [
1 2 200
2 2 300
1 1 400
]
```

Output 7.4.2:

```
trials =
     1     2    200
     2     2    300
     1     1    400
```

You could then determine the duration (column 3) for the second trial (row 2) as `trials(2,3)`. However, you could make the code more transparent using struct. If you did so, you could represent side and brightness with informative names rather than arbitrary numeric codes.

The next example shows a way of representing the data about trials, treating the data as a struct. Each variable assigned in the array `trial` has a numeric index (1, 2, or 3) to designate the trial number. It then has a period and the name of the field: `side`,

brightness, or duration. Having assigned values for the `side`, `brightness`, and `duration` of the stimulus to be shown in trial 1, 2, and 3, you can query the system about the `trial` structure as a whole, about `trial(3)` in particular, about the `side` of trial 2, and about the durations of the stimuli in all trials.

Code 7.4.3:

```
%Initialize struct fields and values
trial(1).side = 'left';
trial(1).brightness = 'bright';
trial(1).duration = 200;

trial(2).side = 'right';
trial(2).brightness = 'bright';
trial(2).duration = 200;

trial(3).side = 'left';
trial(3).brightness = 'dim';
trial(3).duration = 400;

%Examine struct values
trial
t3 = trial(3)
t2_side = trial(2).side
t_durations = [trial(:).duration]
```

Output 7.4.3:

```
trial =
1x3 struct array with fields:
    side
  brightness
   duration
t3 =
         side: 'left'
  brightness: 'dim'
    duration: 400
t2_side =
right
t_durations =
    200    200    400
```

The fields of structures need not be restricted to single values, though in the trial example above, each field (each attribute of a trial) had a single numeric or string value. A field can accommodate a matrix of arbitrary size, as demonstrated below using dependent variables recorded in an experimental session. The structure is called `subject`. It contains, so far, data from two subjects. Subject 1 has reaction times (RTs) and errors for three trials in each of two sessions, whereas subject 2 has RTs and errors for three trials in each of *three* sessions. The fact that the number of sessions is not the same for the two subjects

causes no problems, though it would if you were using a standard matrix. Nor does it cause problems that `subject(2)` has two fields not found in `subject(1)`, namely, `debrief` and `comment`. Note finally that the `comment` field is a `string`, whereas the other fields are numbers. Like cells, structures can accommodate a diversity of types and sizes of elements, making structures, like cells, very useful.

Code 7.4.4:

```
subject(1).RTs = [
    500 400 350
    450 375 325
];
subject(1).errors = [
    10 8 6
    4 3 2
];
subject(2).RTs = [
    600 500 450
    550 475 425
    500 425 400
];
subject(2).errors = [
    10 8 6
    4 3 2
    3 2 1
];
subject(2).debrief = true;
subject(2).comment = 'That was a really cool experiment!';

subject

s1 = subject(1)
s2 = subject(2)
```

Output 7.4.4:

```
subject =
1x2 struct array with fields:
    RTs
    errors
    debrief
    comment
s1 =
    RTs: [2x3 double]
    errors: [2x3 double]
    debrief: []
    comment: []
s2 =
    RTs: [3x3 double]
```

```
        errors: [3x3 double]
        debrief: 1
        comment: [1x34 char]
```

With this structure you can easily write a program to save the mean RT and number of errors made by each subject. The results go to a `.txt` file, named `RTdata.txt`.

Code 7.4.5:

```
outfile = 'RTdata.txt';
outfile = fopen(outfile, 'wt');
% print header line
fprintf(outfile, 'sub\tRT\tErrors\n');
% print data table
for subjectnumber = 1:2
    fprintf(outfile, '%3d\t%5.1f\t%3.1f\n', subjectnumber, ...
        mean(subject(subjectnumber).RTs(:)), ...
        mean(subject(subjectnumber).errors(:)));
end
type('RTdata.txt')
```

Output 7.4.5:

sub	RT	Errors
1	400.0	5.5
2	480.6	4.3

Converting between cell arrays or matrices and structures may seem tedious, but we recommend using structures to keep the representation of data organized and transparent. Happily, there is a shortcut for initializing the elements of a structure without having to write a whole slew of assignment statements.

The `deal` command is used in assignment operations when a different element is to be assigned to the same field of each element of the structure, or a single constant value is to be assigned to each instance of a field of the structure. Note the brackets around the expression to the left of the equals sign in the assignment statements that use `deal`. The first operation initializes `mystruct` as an 8×1 struct array. After that the index `(1:8)` is not needed, as long as your intention is to process all the elements of `mystruct`. The following operations give each element an empty field, a field with a random integer, and a field with an integer that counts down from 8 to 1.

Code 7.4.6:

```
% dealexample.m
clear
clc
% Initializing fields of a struct
```

```
[mystruct(1:8).initiallyZeroVariable] = deal(0);
[mystruct.initiallyEmpty] = deal([]);
[mystruct.random] = ...
    deal(randi(10),randi(10),randi(10),randi(10),...
        randi(10),randi(10),randi(10),randi(10));
[mystruct.integers] = deal(8,7,6,5,4,3,2,1);
ms_1 = mystruct(1)
ms_2 = mystruct(2)
ms_8 = mystruct(8)
```

Output 7.4.6:

```
ms_1 =
    initiallyZeroVariable: 0
           initiallyEmpty: []
                random: 9
           integers: 8
ms_2 =
    initiallyZeroVariable: 0
           initiallyEmpty: []
                random: 7
           integers: 7
ms_8 =
    initiallyZeroVariable: 0
           initiallyEmpty: []
                random: 4
           integers: 1
```

Extracting values from the structure can be accomplished using `deal`, again. These two operations demonstrate conversion of the structure contents to a cell array and then to a numerical matrix when the numerical matrix is the most convenient representation of the data for subsequent computation.

Code 7.4.7:

```
% Reading field of struct to cell array using deal
[TheIntegersCellArray{1:length(mystruct)}] = ...
    deal(mystruct(:).integers)
% Converting cell array to matrix using cell2mat
TheIntegerMatrix = cell2mat(TheIntegersCellArray)
```

Output 7.4.7:

```
TheIntegersCellArray =
    [8]    [7]    [6]    [5]    [4]    [3]    [2]    [1]
TheIntegerMatrix =
     8     7     6     5     4     3     2     1
```


Finally, structures are useful for accessing directories. The contents of the current directory can be represented as a variable of type `struct` if you assign the output of the `dir` command to a variable. If you had an automated multi-step analysis program, you could exploit this feature of `dir` to automatically scan the directory for entries whose names indicate that they contain intermediate data requiring subsequent analysis. The example lists the directory (as it existed when this example was written), then assigns the directory contents to the variable `mydir`. It then uses the `strfind` command, which is described in the next section, to identify a file name that has `'step2.mat'` in it, for further processing,

Code 7.4.8:

```
dir
mydir = dir
fprintf('\nFiles found:\n');
for i = 1:length(mydir)
    fname = mydir(i).name;
    if strfind(fname,'step2.mat')
        fprintf(['File named '%s'' ...
                'will be processed for step 3.\n'],fname);
    end
end
end
```

Output 7.4.8:

```
.                my_dlm_data.txt
..               mydata.txt
Apps             mydata1.txt
DatafromStep1.mat myexpt_EF_062913_step2.mat
Days_In_A_Month.m precisionexample.m
DirectoryExample.m readRTdata.m
SimonDemo.m      rtdata.txt
SimonDemo2.m     sampledata.txt
SimonDemo3.m     simondata.mat
booleanloopexample.m syncme.m
daysinMonth.m
garbage.m

mydir =
22x1 struct array with fields:
    name
    date
    bytes
    isdir
    datenum

Files found:
File named 'myexpt_EF_062913_step2.mat' will be processed
for step 3.
```

The `date` and `bytes` fields of the directory struct array are available if you need to refer to creation date or file size. `mydir(i).isdir` will be 1 if the *i*-th entry is itself a directory (a sub-folder). Otherwise, `mydir(i).isdir` will be 0 to indicate the entry is a file. The `datenum` field records the creation date in numerical form.

Structures can be organized hierarchically to clarify the organization of data, as in the following example, which represents the reaction time on the fourth trial of the third block of the second subject in an experiment, and the second trial of the fourth block of the fifth subject. As you will see, the fields of a structure can be structures. The data can be scanned with nested `for` loops, as shown in the `Analysis loop` section (whose output is not shown).

Code 7.4.9:

```
thisRT = subject(2).block(3).trials(4).RT;
thisRT = subject(5).block(4).trials(2).RT;

% Analysis loop
for subcount = 1:5
    for blockcount = 1:4
        for trials = 1:10
            thisRT = ...
subject(subcount).block(blockcount).trials(trialcount).RT;
            % further analysis of thisRT...
        end
    end
    % Output for this subject would go here
end
```

7.5 Searching and Modifying Strings

Earlier in this book, you were exposed to ways of comparing and manipulating numbers in matrices. MATLAB provides ways of performing the same kinds of manipulations on strings.

The most common need for a string comparison is to test two strings for equality, or to determine if a particular substring is embedded within a longer string. Using the familiar double equals sign (`==`) to compare two strings would seem to solve this problem, but if the strings were of unequal length, the result would not be what you hoped for. Your program would halt and you would get an error message.

The problem is solved with a function called `strcmp`, short, presumably, for string comparison. This function takes two string arguments and returns a 1 or 0, depending on the identity of the two strings, and reports the strings as “different” (i.e. returns 0) when they are not the same length, rather than halting with an error message. A variant of `strcmp`, called `strcmpi`, performs the same comparison on two strings, while ignoring case differences.

Code 7.5.1:

```
'apples' == 'oranges'
apples_to_oranges = strcmp('apples', 'oranges')
apples_to_apples = strcmp('apples', 'apples')
apples_to_APPLES = strcmp('apples', 'APPLES')
apples_to_APPLES_ignoring_case = ...
    strcmpi('apples', 'APPLES')
```

Output 7.5.1:

```
Error using ==
Matrix dimensions must agree.

apples_to_oranges =
     0
apples_to_apples =
     1
apples_to_APPLES =
     0
apples_to_APPLES_ignoring_case =
     1
```

The `strfind` command, which you encountered in Code 7.4.8, takes two string arguments and reports the character position where any instance of the second string is embedded in the first. If there is more than one match, `strfind` returns a matrix of the letter positions in the longer string where each of the instances appears. If there are no such instances, `strfind` returns an empty string. If you are interested in simply detecting one or more instances of a target substring, you can use the `any` operator on the result returned from `strfind`.

Code 7.5.2:

```
s = ['How much wood could a wood chuck chuck'...
     ' if a wood chuck could chuck wood?'];
all_wood_in_s = strfind(s, 'wood')
all_could_in_s = strfind(s, 'could')
all_should_in_s = strfind(s, 'should')
any_wood_in_s = any(strfind(s, 'wood'))
any_should_in_s = any(strfind(s, 'should'))
```

Output 7.5.2:

```
all_wood_in_s =
     10     23     45     68
all_could_in_s =
     15     56
all_should_in_s =
```

```

    []
any_wood_in_s =
    1
any_should_in_s =
    0

```

Another useful function is `strrep`, which stands for “string replacement.” `strrep` takes three string arguments. The first is the string to be modified. The second is the substring to be found in that string. The third is a new substring that will replace each instance of the second substring. The result is a new (perhaps modified) string that can be assigned to a variable.

Code 7.5.3:

```

s = ['How much wood could a wood chuck chuck'...
     ' if a wood chuck could chuck wood?'];
s1 = strrep(s, 'wood', 'cider');
s2 = strrep(s1, 'chuck', 'press');
s2

```

Output 7.5.3:

```

s2 =
How much cider could a cider press press if a cider press
could press cider?

```

You might find use for the `strrep` operation to modify a file name if you were reading a data file from an experiment and wanted to use a variant of the same name for the output file, or if you were reading a `.mat` file (see Section 6.14) and generating the results in `.txt` format for use by another program.

Code 7.5.4:

```

infilename = 'myexpt_EF_062913_step2.mat'
outfilename = strrep(infilename, 'step2', 'step3')
fprintf('\n');
inMatName = 'MeanReactionTimes.mat'
outTxtName = strrep(inMatName, '.mat', '.txt')

```

Output 7.5.4:

```

infilename =
myexpt_EF_062913_step2.mat
outfilename =
myexpt_EF_062913_step3.mat

inMatName =
MeanReactionTimes.mat

```

```
outTxtName =  
MeanReactionTimes.txt
```

In reading data files you may often find that you have to read both strings and numbers. Consider the text file, `RTdata.txt`, created in Code 7.4.5. The first line of the file is the header line, which must be read as text, whereas subsequent lines are numbers representing the subject, RTs, and errors, which must be read as numbers. You can first read each line into a string variable (`headers` or `nextline`) and then read that string using `textscan`. For headers, you can read into a cell array using a string format (`%s`), and for `nextline` you can read into an array using a numeric format (`%f`). To keep reading until the end of the file, the reading of lines after the header is embedded in a `while ~feof(infile)` loop, which repeats until the reading of the last line of the file is signaled by the function `feof` returning a value of 1, indicating “found end of file.” Finally, you can print the headers and numbers in a convenient format by transposing the output of `textscan` from columns to rows, using the transpose (`'`) operator.

Code 7.5.5:

```
infilename = 'RTdata.txt';  
infile = fopen(infilename);  
firstline = fgetl(infile); %read the header line  
headers = textscan(firstline, '%s');  
cell_of_headers = headers{1}(1:3)'  
  
matrix_of_numbers = [];  
while ~feof(infile)  
    nextline = fgetl(infile);  
    nextvalues = textscan(nextline, '%f');  
    matrix_of_numbers =...  
        [matrix_of_numbers; nextvalues{1}(1:3)'];  
end  
  
fclose(infile);  
matrix_of_numbers
```

Output 7.5.5:

```
cell_of_headers =  
    'sub'    'RT'    'Errors'  
matrix_of_numbers =  
    1.0000    400.0000    5.5000  
    2.0000    480.6000    4.3000
```

7.6 Applying Data Types

Suppose you have data in a tab-delimited file named `Simon.txt`, and you wanted to compute the mean reaction time for the correct trials. See Output 7.6.1. The code used to generate it is not shown here.

Output 7.6.1:

Trial	side	stim	comp	Key	Resp.	RT
1	L	L	C	R	error	0.73
2	R	R	C	R	correct	0.79
3	L	R	I	R	correct	0.54
4	R	L	I	L	correct	0.51
5	L	R	I	R	correct	0.44
6	L	L	C	L	correct	0.49
7	R	L	I	R	error	0.39
8	R	R	C	R	correct	0.68

...data from many more trials not shown

Each line of this file (after the first) has seven variables: a number, four characters, a string, then a number. The following code will parse the lines, extract the variables of different types, and compute the mean.

Code 7.6.2:

```
% DoSimon.m
fin = fopen('Simon.txt');
allRTs = [];
%Skip the header line
headerline = fgetl(fin);
while ~feof(fin)
    aline = fgetl(fin);

    %Read in the variables
    cellvalues = textscan(aline, '%d %s %s %s %s %s %f');
    Trnum = cell2mat(cellvalues(1));
    side = char(cellvalues{2});
    stim = char(cellvalues{3});
    comp = char(cellvalues{4});
    Key = char(cellvalues{5});
    Resp = char(cellvalues{6});
    RT = cell2mat(cellvalues(7));

    % Assemble the correct trial RT's
    if strcmp(Resp, 'correct')
        allRTs = [allRTs RT];
    end
end
meanRT = mean(allRTs);
fprintf('Mean of correct RTs is %f\n', meanRT);
```

Output 7.6.2:

```
| Mean of correct RTs is 0.575000
```

The `txtscan` command reads each line's seven variables into a 1×7 cell array, `cell-values`, which is a mixture of integers, real numbers, characters, and strings. Before you can do further analysis, each cell must be translated to the corresponding standard MATLAB variable type. Take special note of the differences between parentheses and braces in the assignment statements that make this translation for each of the variables. They are tricky!

7.7 Practicing Data Types

Try your hand at the following exercises, using only the methods introduced so far in this book or in information given in the problems themselves.

Problem 7.7.1:

Create a 5×3 cell array, `G`, with students' names (Adam, Brad, Charley, David, or Emily) in the first column of the cell array; each student's corresponding numerical average (90, 92, 96, 95, 88) in the second column of the cell array; and each student's letter grade (A-, A-, A, A, B+) in the third column.

Represent the same data as above in a 5×1 struct array, `studentStruct(1:5)`, with two fields, `name`, and `average` initialized as above. Write a program to compute the letter grade based on `studentStruct(i).average` and record it in `studentStruct(i).letter` for each student.

Problem 7.7.2:

Create a cell array, `C`, whose rows 32 through 127 contain that number as an integer in the first column and the character equivalent of that number in the second column.

Problem 7.7.3:

Use `fprintf` to print the numbers 65 through 90 in one column, the character equivalent of that value in column 2, the numbers 97 through 122 in column 3, and the character equivalent of that number in column 4.

Problem 7.7.4:

Write a program to administer a computerized questionnaire on a topic of interest to you. Use a structure data type and allow participants to answer with whole sentences or phrases for at least some items. Save the data in an external file. Record the time to answer each question.

Problem 7.7.5:

Generate a data set using Code 7.7.5 and verify the accuracy of your program by comparing its checksum output with that in Output 7.7.5.

Code 7.7.5:

```
% Code_7_6_5.m
rng('default')
for n = 1:20
    r1 = randn;
    r2 = mean([r1 r1 randn]) + .4;
    subject(n).score1 = r1;
    subject(n).score2 = r2;
end
subject
checksum1 = sum([subject(:).score1])
checksum2 = sum([subject(:).score2])
```

Output 7.7.5:

```
subject =
1x20 struct array with fields:
    score1
    score2
checksum1 =
    4.5867
checksum2 =
   13.5765
```

Now, compute the correlation coefficient between `score1` and `score2`.

Problem 7.7.6:

Evaluate the difference between `score1` and `score2` using a correlated (within-subjects, or paired-samples) t-test. As a reminder, the computational formula for a within-subjects t-test is given below, where $d = \text{score2} - \text{score1}$ (for each subject), and n is the number of subjects, 20. If the absolute value of your computed value of t is greater than 2.861, the difference between `score1` and `score2` is statistically significant at the .01 level ($p < .01$, two-tailed test, with $df = 19$). Verify the accuracy of your computation and decision using the statistical package of your choice.

$$t = \frac{\sum d}{\sqrt{\frac{n(\sum d^2) - (\sum d)^2}{n-1}}}$$

8. Modules and Functions

This chapter covers the following topics:

- 8.1 Taking a top-down approach to programming by using modules
- 8.2 Writing and using general-purpose functions
- 8.3 Getting multiple outputs from functions
- 8.4 Passing multiple input arguments to functions
- 8.5 Creating multiple functions in a file
- 8.6 Calling functions properly
- 8.7 Exploiting recursive functions
- 8.8 Drawing on previously defined functions versus creating your own
- 8.9 Practicing modules and functions

The commands that are introduced and the sections in which they are premiered are:

<code>end (function)</code>	(8.2)
<code>return</code>	(8.2)
<code>function</code>	(8.2)

8.1 Taking a Top-Down Approach to Programming by Using Modules

All the programs presented so far are relatively small because they merely illustrate different approaches to larger programming needs. As programs grow, they tend to become more complex, but with greater program length and complexity, programs can get hard to follow, leaving you feeling like a rat lost in a maze.

The purpose of this chapter is to prevent such “rats’ nests.” Expressed more positively, the aim of the chapter is to help you create code that is clear and flexible. Code can be clear if it is designed in a modular fashion (i.e., broken into meaningful sub-programs). It can be flexible if it is equipped with general-purpose functions. The next several sections focus on functions. The present section focuses on modules. The latter term is one we use to refer to stand-alone scripts that perform one or a small number of instructions. The term “modules” is not an official MATLAB term.

To illustrate the value of modular programming, consider the following example, which is a program for selecting students for admission to a college. Here is a script (saved to the file `College_Admissions_5.m`) that illustrates how the selection procedure might work. Rest assured that this program is not actually being used at any institution of higher education, at least as far as we know. The code is less transparent than it might be by design. Just skim it because a simpler, more modular, version will follow.

Code 8.1.1:

```
% College_Admissions_5

% Assuming that SATs and GPAs are related to IQs,
% this program generates dummy data for SATs, GPAs,
% Extra- curriculars (EC), and distance (Dist) from the
% college, giving larger scores to greater distance from
% the college (for geographical diversity).
% The SATS and GPAs are summed, each of the student's
% three new scores (Acad, EC, and Dist) are normed, and
% then the min required score for admission is gradually
% increased until the number admitted no longer exceeds
% max_admits_allowed.

% Clear variables, clear and open the commandwindow
clear all
clc
commandwindow

% Set constants
applications = 30;
max_admits_allowed = 10;

IQmean = 110;
IQsd = 20;
SATQmean = 500;
SATQsd = 100;
SATVmean = 500;
SATVsd = 100;
ECsd = 10;
GPAmean = 2.0;
GPAsd = 10;

% Preallocate arrays using deal
[IQ SATQ SATV GPA Acad EC Dist] = deal(zeros(applications,1));

% Generate dummy scores to test the program
IQ = IQmean + (randn(applications,1)) * IQsd;
SATQ = SATQmean + (randn(applications,1)) * SATQsd;
SATV = SATVmean + (randn(applications,1)) * SATVsd;
GPA = GPAmean + (randn(applications,1)) * GPAsd;
EC = abs(randn(applications,1)) * ECsd;
Dist = abs(randn(applications,1));
Acad = SATQ + SATV + 100 * GPA;

% Normalize the scores
Acad = (Acad - min(Acad)) ./ (max(Acad)-min(Acad));
```

```
EC = (EC -min(EC)) ./ (max(EC)-min(EC));
Dist = (Dist - min(Dist)) ./ (max(Dist)-min(Dist));

% Create a Scores matrix, including, in the final column,
% each student's total score
Scores = [[1:applications]' Acad EC Dist];
Scores(:,5) = [Acad + EC + Dist];

% Admit the top max_admits_allowed students (plus any ties)
SortedScores = sortrows(Scores,-5);
criterion = SortedScores(max_admits_allowed,5);
SortedScores(:,6) = 0;
SortedScores((SortedScores(:,5) >= criterion),6) = 1;
ScoresAndAcceptances = sortrows(SortedScores,1);

% Display the results
fprintf('App.\tAcad.\tExtra.\tDist.\tTotal\tAccept\n\n')
fprintf('%4d\t%6.2f\t%6.2f\t%6.2f\t%6.2f\t%4d\n', ...
        ScoresAndAcceptances)
fprintf('\r')
Students_Accepted = find(ScoresAndAcceptances(:,6));
fprintf('Accepted Students:\n');
fprintf('%3d', Students_Accepted);
fprintf('\n\n')
fprintf('Cutoff score: %5.03f\n', criterion);
```

Output 8.1.1:

App.	Acad.	Extra.	Dist.	Total	Accept
1	0.48	0.26	0.34	1.08	0
2	0.33	0.03	0.45	0.82	0
3	0.21	0.74	0.08	1.03	0
4	0.24	0.91	0.09	1.24	0
5	0.70	0.17	0.30	1.17	0
6	0.59	0.00	0.09	0.68	0
7	0.79	0.41	0.10	1.30	0
8	0.53	0.24	0.33	1.10	0
9	0.56	0.08	0.43	1.07	0
10	0.11	0.07	0.71	0.89	0
11	0.43	0.31	0.31	1.05	0
12	0.67	0.92	0.36	1.96	1
13	0.51	0.05	0.00	0.56	0
14	0.98	0.48	0.01	1.48	1
15	0.00	0.08	0.02	0.10	0
16	0.52	1.00	0.22	1.75	1
17	0.64	0.00	0.07	0.71	0
18	0.41	0.61	0.10	1.12	0
19	0.62	0.44	0.49	1.56	1

20	1.00	0.51	0.42	1.93	1
21	0.34	0.67	0.06	1.08	0
22	0.51	0.10	0.38	1.00	0
23	0.19	0.11	0.14	0.44	0
24	0.82	0.24	0.39	1.46	1
25	0.21	0.04	0.28	0.53	0
26	0.38	0.43	1.00	1.81	1
27	0.98	0.49	0.17	1.64	1
28	0.83	0.28	0.57	1.68	1
29	0.41	0.41	0.17	0.99	0
30	0.92	0.23	0.28	1.42	1

Accepted Students:

12 14 16 19 20 24 26 27 28 30

Cutoff score: 1.423

Code 8.1.1 may be hard to follow because it is lengthy and intricate. The program was written with an outline in mind, but the outline is not readily apparent in the code.

The code below shows how the same material can be organized as a series of distinct scripts, or “modules.” Organizing the code in a modular fashion reflects a top-down approach to programming rather than a bottom-up approach. It is useful to take a top-down as well as a bottom-up approach to programming because the top-down approach helps you focus on large-scale organization. When you are working at a more detailed level, within a module, you can concentrate on the minutia that, unavoidably, must be considered. An entirely bottom-up approach, by contrast, forces you to focus on the syntax of individual lines of code. Generating code in a top-down fashion becomes more natural as the lower-level details become more automatic. This is why modules and functions are introduced at this point in the book rather than earlier.

In the material that follows, Code 8.1.1 has been broken down into modules (Codes 8.1.2 through Code 8.1.9), each of which was previously stored as a stand-alone .m-file script (see Chapter 2). Each module is in its own file, and can be called from another module, in just the same way a program can be called from the Command window. Code 8.1.2 is the main program, Code 8.1.3 is the first module called by the main program, Code 8.1.4 is the second module called by the main program, and so on. Each called program indicates, via a comment, which program called it (the main program in this case). Commented references to calling programs help you keep track of the lineage of your code.

Code 8.1.2:

```
% College_Admissions_Main.m

Clear_Start;
Set_Constants;
Generate_Dummy_Scores;
Normalize_Scores;
Create_Scores_Matrix;
Select_Students;
Display_Results;
```

Code 8.1.3:

```
% Clear_Start.m
% Called by College_Admissions_Main.m
clear all
clc
commandwindow
```

Code 8.1.4:

```
% Set_Constants.m
% Called by College_Admissions_Main.m
applications = 30;
max_admits_allowed = 10;
IQmean = 110;
IQsd = 20;
SATQmean = 500;
SATQsd = 100;
SATVmean = 500;
SATVsd = 100;
ECsd = 10;
GPAmean = 2.0;
GPAsd = 10;
[IQ SATQ SATV GPA Acad EC Dist] = deal(zeros(applications,1));
```

Code 8.1.5:

```
% Generate_Dummy_Scores.m
% Called by College_Admissions_Main.m
IQ = IQmean + (randn(applications,1)) * IQsd;
SATQ = SATQmean + (randn(applications,1)) * SATQsd;
SATV = SATVmean + (randn(applications,1)) * SATVsd;
GPA = GPAmean + (randn(applications,1)) * GPAsd;
EC = abs(randn(applications,1) * ECsd);
Dist = abs(randn(applications,1));
Acad = SATQ + SATV + 100 * GPA;
```

Code 8.1.6:

```
% Normalize_Scores.m
% Called by College_Admissions_Main.m
Acad = (Acad - min(Acad)) ./ (max(Acad) - min(Acad));
EC = (EC - min(EC)) ./ (max(EC) - min(EC));
Dist = (Dist - min(Dist)) ./ (max(Dist) - min(Dist));
```

Code 8.1.7:

```
% Create_Scores_Matrix.m
% Called by College_Admissions_Main.m
Scores = [[1:applications]' Acad EC Dist];
Scores(:,5) = Acad + EC + Dist;
```

Code 8.1.8:

```
% Select_Students.m
% Called by College_Admissions_Main.m
SortedScores = sortrows(Scores,-5);
criterion = SortedScores(max_admits_allowed,5);
SortedScores(:,6) = 0;
SortedScores((SortedScores(:,5) >= criterion),6) = 1;
ScoresAndAcceptances = sortrows(SortedScores,1);
```

Code 8.1.9:

```
% Display_Results.m
% Called by College_Admissions_Main.m
fprintf('App.\tAcad.\tExtra.\tDist.\tTotal\tAccept\n\n')
fprintf(...
    '%4d\t%6.2f\t%6.2f\t%6.2f\t%6.2f\t%4d\n', ...
    ScoresAndAcceptances)
fprintf('\r')
Students_Accepted = find(ScoresAndAcceptances(:,6));
fprintf('Accepted Students:\n');
fprintf('%3d', Students_Accepted);
fprintf('\n')
fprintf('Cutoff score: %5.03f\n', criterion);
```

Of all the programs listed above (Codes 8.1.2 through 8.1.9), only one needs to be run directly by the user: the main program, `College_Admissions_Main` (Code 8.1.2). When that program is run, it calls each of the programs listed within it, one after the other. When each of those programs finishes, it automatically returns control to the program that called it. The output is the same as before (Output 8.1.1).

One other feature of modular programming that makes the approach appealing is that when you have multiple files open in the editor, you can easily switch from one to the other by clicking on one of the filenames listed in the tab buttons in the Editor window.

8.2 Writing and Using General-Purpose Functions

A reason why the programs in Codes 8.1.2 through 8.1.9 work is that they make use of the same variables. Thus, `Generate_Dummy_Scores.m` (Code 8.1.5) makes use of the

values created in `Set_Constants.m` (Code 8.1.4). Similarly, `Normalize_Scores.m` (Code 8.1.6) makes use of the values created in `Generate_Dummy_Scores.m` (Code 8.1.5). The reason the variables from all the modules are universally accessible is that all the programs use the same workspace.

Having all modules use a common workspace can be a great convenience. On the other hand, there are times when this can be a nuisance. Those are the times when functions are used. What are functions in MATLAB, and why does a common workspace tend to be a nuisance? Are functions only nuisances, or do they have redeeming qualities? The answers to the last two questions, it turns out, are, resoundingly, No and Yes, respectively. In other words, functions are good! Here's why.

Functions in MATLAB are basically the same as ordinary functions in mathematics. They take inputs, and they generate outputs. The relation between the input of a function and the output of a function, whether in math or in MATLAB, is what defines the function. Going from the input to the output is what the function does.

Functions have two important assets for programming. One is generality. When a function is used, it generates an output from any acceptable input. The second asset of functions is that they effectively hide the complexities of the computations they employ, which can be distracting if you are working (trying to think) at a higher level.

Though we are introducing functions here explicitly, you have actually been introduced to them many times in this book. This happened when you were exposed to function calls, as in `mean`, `median`, `disp`, and `double`. These are functions built in to MATLAB.

How do you write your own functions? Toward answering this question, recall the syntax for a function call. Consider this simple example.

Code 8.2.1:

```
| r = [1:99];  
| mean_r = mean(r)
```

Output 8.2.1:

```
| mean_r =  
|      50
```

When the `mean` function is called, it computes the arithmetic average of `r`, taking the values of `r` as input for the necessary calculations.

Here are some examples of computing the mean of several sets of values that do not take advantage of the built in `mean` function:

Code 8.2.2:

```
| meanA = (1+3+5+7+9)/5  
| a = pi;  
| b = 1492;
```

```
c = 6.02;
meanB = (a+b+c)/3
meanC = sum(1:10)/10
```

Output 8.2.2:

```
meanA =
      5
meanB =
 500.3872
meanC =
  5.5000
```

The input to a function is sometimes referred to as the *argument* for the function. When a function is called in MATLAB, it assigns the argument to the function as input. The function then returns output to the calling program.

Let's now write a new function, `mymean.m`, that will compute the means for the cases above. We avoid using the filename `mean.m` because our new function would replace the built-in `mean` function, which is not a good idea!

Code 8.2.3:

```
% function mymean.m
function myresult = mymean(inputarray);
myresult = sum(inputarray)/length(inputarray)
return
```

Notice that the function ends with `return`. This term is optional, but it is helpful to include it to indicate where the function concludes.

Once this function has been defined in its own file (`mymean.m`), it can be called from another module or function as in the three calls below.

Code 8.2.4:

```
meanD = mymean([1 3 5 7 9])
meanE = mymean([pi 1492 6.02])
meanF = mymean([1:10])
```

Output 8.2.4:

```
meanD =
      5
meanE =
 500.3872
meanF =
  5.5000
```


A more detailed example of function use follows. Here we introduce a new function called `normalize` which we create after realizing that it would be useful to have a general-purpose function to translate the values in any given numerical array to values ranging from 0 to 1, where 0 is assigned to the smallest value, 1 is assigned to the largest value, and values in between are assigned values corresponding to their distance from the minimum, divided by the distance of the maximum from the minimum (see Code 8.1.6). It would be useful to create such a function because it would be inconvenient to have to change the variable names in Code 8.1.6 to some other set of names for every other normalizing problem. Similarly, it would be confusing to stick with the names originally used (e.g., `Acad`) in some other context where `Acad` is not relevant (e.g., amusement park ratings). We want a function that carries out computations on variables with generic names, such as `x` and `y` that are meaningful only within the called function while the calling program can use different meaningful names, such as `AcademicRank` or `AmusementParkRating`.

There are several points to keep in mind about functions. First, a function must be saved as a `.m` script or it must be included in a `.m` file that defines a function. Second, the name of the saved `.m` script can be used to call the function. Third, within the `.m` file itself, the first term of the first executable line (after any comments) must be the word `function`. Fourth, the syntax of the first executable line of every function must be of the following form, where `input` denotes the function's argument (it needn't be called `input`) and `output` denotes the function's result (it needn't be called `output`).

```
function output = name_of_function(input)
```

Fifth, the subsequent executable line or lines of code constitute the operations to be performed until the end of the function is reached, as indicated by a `return` or `end` statement, the end of the file, or a new function definition.

Here is code for the new function, `normalize`, which takes one input argument (a matrix `x`) and returns an output argument, `y`, with the same size as `x`, representing the values of `x` normalized to a range 0 through 1.

Code 8.2.4:

```
% normalize.m
function y = normalize(x)
y = (x - min(x)) ./ (max(x)-min(x));
end
```

We can check that the new function works.

Code 8.2.5:

```
x = [1:8]
normalized_values = normalize(x)
```

Output 8.2.5:

```
x =
     1     2     3     4     5     6     7     8
normalized_values =
```

	0	0.1429	0.2857	0.4286	0.5714	0.7143
0.8571	1.0000					

Note that in Code 8.2.5, the name of the input array passed as an argument to `normalize` is `x`. `x` is also the name of the variable used in `normalize`. Will the function still work if the name of the argument isn't the same as the name of the variable used in the function? The following example shows that it will, demonstrating that the function takes the argument supplied to the function by the calling program (in this next case, the array called `a`) and substitutes it for its own input variable (in this case, the array called `x`) in all computations in the function.

Code 8.2.6:

```
a = [1:8];
normalized_values = normalize(a)
```

Output 8.2.6:

x =							
	1	2	3	4	5	6	7
normalized_values =							8
	0	0.1429	0.2857	0.4286	0.5714	0.7143	
0.8571	1.0000						

What happens if we ask for the value of `y`, which is the name of the output generated in `normalize` (see Code 8.2.4), after `normalize` has returned its output and we are back in the calling program?

Code 8.2.7:

```
a = [1:8];
normalize(a);
y
```

Output 8.2.7:

```
??? Undefined function or variable 'y'.
```

Surprisingly, we get an error message. MATLAB tells us that `y` is an undefined function or variable. What did we do wrong?

Nothing! The reason for the message is that variables inside functions are *local* variables, not *global* variables. Local variables are restricted to the variable workspace that is exclusively reserved for the function. The designers of MATLAB appreciated that much as one might want to use special, generic terms inside a variety of functions (e.g., `x` in a function that normalizes, `x` in a function that returns the mean, and so on), it would be best to keep the variables inside functions restricted to, or “local” to those functions, at least by default.

8.3 Getting Multiple Outputs From Functions

The function `normalize` generates only one output. However, MATLAB functions can give multiple outputs. Consider this example, a function that calculates a median split—so it finds values above and below the median—and then normalizes the scores in the lower half separately from the scores in the upper half.

Code 8.3.1:

```
function [ly, uy] = normalize_split(x)

lx = x(x <= median(x));
ux = x(x > median(x));
uy = (ux - min(ux)) ./ (max(ux) - min(ux));
ly = (lx - min(lx)) ./ (max(lx) - min(lx));
return
```

We can check that the function works by calling it. In so doing, we must be sure that each of the two outputs, `uy` and `ly`, are mapped to variables available to the calling program. In this case, we refer to the mapped output variables as `lower_normed` and `upper_normed`, respectively.

Code 8.3.2:

```
a = randi(10,1,14);
[lower_normed, upper_normed] = normalize_split(a)
```

Output 8.3.2:

```
lower_normed =
    0.3333    0.3333    0.3333    0.3333         0         0
    1.0000
upper_normed =
    1.0000    1.0000    0.2500    0.2500    1.0000         0
    0.5000
```

Note that the outputs, like the inputs, can have different names in the function and in the calling program. Typically, the name in the calling program will be specific to the problem the main program addresses, whereas the variable name in the function can be generic, demonstrating the abstract utility of the function.

Whether or not a function runs “to the very end” is optional. It may be that for some conditions, the function should return after some operations have been performed. Here is an example of a function to return the square root, after checking to see if the argument is positive, assuming, in this case, that square roots of negative numbers are not allowed because imaginary numbers fall outside the acceptable purview. The `return` at the end is optional, but reminds the programmer of the program flow.

Code 8.3.3:

```
% myroot.m
function result = myroot(x)
if x > 0
    result = sqrt(x);
else
    fprintf(...
        '\nCan''t take the square root of a negative number\n');
    result = NaN;
end
return
```

Code 8.3.4:

```
Result_1 = myroot(2)
Result_2 = myroot(-2)
```

Output 8.3.4:

```
Result_1 =
    1.414213562373095

Can't take the square root of a negative number
Result_2 =
    NaN
```

8.4 Passing Multiple Input Arguments to Functions

Calls to functions can have more than one input argument. When multiple input arguments are supplied to a function, they are assigned in the order in which they are specified in the function call (from left to right). It is usually required that the number of variables in the call to the function match the number of input variables in the function definition. (Some functions can specify a variable number of arguments, which we don't address here).

Here is an example in which `normalize_split_two_args` takes two arguments rather than one, contrary to the previous examples. The first argument is an $n \times 1$ numeric array. The second argument is of type `string`. The function does different things depending on the second argument. It normalizes scores above or below the *median* of the input array if the second argument is `median`, but it normalizes scores above or below the *mean* of the input array if the second argument is `mean`. If the second argument is neither `median` nor `mean`, an error message is shown.

Code 8.4.1:

```
% normalize_split_two_args.m
% Splits array in first argument into
```

```
% lower and upper halves, using the
% criterion ('mean' or 'median')
% specified in the second argument

function [ly, uy] = normalize_split_two_args(x,typeofsplit);

lx = [];
ux = [];

if strcmp(typeofsplit,'median') % median split
    lx = x(x<=median(x));
    ux = x(x>median(x));

elseif strcmp(typeofsplit,'mean') % mean split
    lx = x(x<=mean(x));
    ux = x(x>mean(x));

else % error feedback
    disp(['Error: An invalid type of split'...
        ' in the call to normalize_split_two_args']);
    [ly, uy] = deal(NaN);
    return
end

ly = (lx - min(lx)) ./ (max(lx) - min(lx));
uy = (ux - min(ux)) ./ (max(ux) - min(ux));
return
```

Calls to `normalize_split_two_args` follow, after which the output is shown. As numerical input to the function, we use a logarithmically spaced array, which has the property that its mean (25.8) and median (10.5) differ.

Code 8.4.2:

```
a = logspace(0,2,8)
[median_based_lower_norm,median_based_upper_norm_mean] = ...
    normalize_split_two_args(a,'mean')
[mean_based_lower_norm,mean_based_upper_norm] = ...
    normalize_split_two_args(a,'median')
[other_based_lower_norm,other_based_upper_norm_mean] = ...
    normalize_split_two_args(a,'anyOtherTerm')
```

Output 8.4.2:

```
a =
    1.0000    1.9307    3.7276    7.1969   13.8950   26.8270
  51.7947  100.0000
median_based_lower_norm =
         0    0.0722    0.2115    0.4806    1.0000
```

```

median_based_upper_norm_mean =
    0    0.3412    1.0000
mean_based_lower_norm =
    0    0.1502    0.4402    1.0000
mean_based_upper_norm =
    0    0.1502    0.4402    1.0000
Error: An invalid type of split in the call to normalize_split
other_based_lower_norm =
    NaN
other_based_upper_norm_mean =
    NaN

```

8.5 Creating Multiple Functions in a File

You can call a function from another function, just as you can call a function from another ordinary program. For example, the functions `min` and `max` (which are built-in MATLAB functions) were called in the function `normalize_split` (Code 8.3.1).

Knowing that a function expressed in one program can call a function expressed in another program may lead you to believe that every function must stand alone. Stand-alone functions, whether provided by MATLAB or written by you, can be useful. However, it is not the case that every function must occupy its own program. More than one function can be fully expressed in the same program. The only proviso is that a function called from inside such a program—a so-called *local* function—cannot be called directly from another program.

Local functions work just like stand-alone functions in that the variables defined in each local function are invisible to other functions. Similarly, all communication between the functions is via the arguments that are passed and returned. The code file containing a local function must itself begin with a function. You can't put local functions in a script that is not a function.

Here is a function that contains a local function. The main function, called `mean_and_trimmed_mean`, would be called by another program or, equivalently, via a command in the Command window. The main function returns the mean of the array as well as the mean of the trimmed array, which trimmed by omitting its largest and smallest values. The array is trimmed via the local function `trimmed`. The local function `trimmed` is invisible to any code outside the file `mean_and_trimmed_mean.m`.

Code 8.5.1:

```

% mean_and_trimmed_mean.m
function [y,ty] = mean_and_trimmed_mean(x)
y = mean(x);
ty = mean(trimmed(x));
return

function zz = trimmed(w)
w = sort(w);

```

```
| zz = [w(2:end-1)];  
| return
```

The call to the function and resulting output follow.

Code 8.5.2:

```
| x = randperm(5).^2  
| [theMean theTrimmedMean] = mean_and_trimmed_mean(x)
```

Output 8.5.2:

```
| x =  
|      25      9      16      4      1  
| theMean =  
|      11  
| theTrimmedMean =  
|      9.6667
```

If you included the code for the function `trimmed` within an ordinary program module—that is, in a `.m` file that does not begin with a function definition—it would not work, as shown here.

Code 8.5.3:

```
| % ComputeMeans_Fails.m  
| x = randperm(5).^2;  
| [theMean theTrimmedMean] = mean_and_trimmed_mean(x)  
  
| function [y,ty] = mean_and_trimmed_mean(x)  
| y = mean(x);  
| ty = mean(trimmed(x));  
| return  
  
| function zz = trimmed(w)  
| w = sort(w);  
| zz = [w(2:end-1)];  
| return
```

Output 8.5.3:

```
| Error: File: ComputeMeans_Fails.m Line: 5 Column: 1  
| Function definitions are not permitted in this context.
```

However, if the program is recast by making it begin with the function `main`, it succeeds. (The `return` and `end` at the end of each function are optional in this case, but useful for clarity). The function `main` is still called by routines outside this file by its filename, `ComputeMeans_Succeeds`.

Code 8.5.4:

```
% ComputeMeans_Succeeds.m
function main
x = randperm(5).^2
[theMean theTrimmedMean] = mean_and_trimmed_mean(x)
return
end % main function

function [y,ty] = mean_and_trimmed_mean(x)
y = mean(x);
ty = mean(trimmed(x));
return
end % mean_and_trimmed_mean

function zz = trimmed(w)
w = sort(w);
zz = [w(2:end-1)];
return
end % trimmed
```

Output 8.5.4:

```
theMean =
    11
theTrimmedMean =
    9.6667
```

“Nested” functions provide a third way of defining functions. A function can be *nested* within a main function, prior to the end statement that ends the main function (when you are using nested functions, every function *must* end with an `end`). The main and nested functions may, of course, intercommunicate in the normal way by passing and returning arguments just as stand-alone and local functions do, but there is an important shortcut, which makes nested functions intrinsically different from main and local functions: If the same variable name is used *both* in the main function and in the nested function, that variable is visible to *both* the main function *and* the nested function. Nested functions are particularly useful for operations that are repeatedly executed but only within the context of a particular function.

Here we use nested functions to carry out the same computations as before. The variables `x`, `theMean`, and `theTrimmedMean` are accessible to all functions without having to be passed as arguments, since these variables are used in both the main function and the nested functions. The MATLAB editor signals that the subfunctions are nested by indenting them when the code is automatically formatted.

Code 8.5.5:

```
% ComputeMeans_Nested.m
function main
x = randperm(5).^2
```



```
mean_and_trimmed_mean
theMean
theTrimmedMean
return

    function mean_and_trimmed_mean
        theMean = mean(x);
        trimmed;
        return
    end % function mean_and_trimmed_mean

    function trimmed
        w = sort(x);
        theTrimmedMean = [w(2:end-1)];
        return
    end % function trimmed

end %function main
```

Output 8.5.5:

```
theMean =
    11
theTrimmedMean =
    9.6667
```

Returning to the example that opened this chapter, Section 8.1 presented two ways of organizing the program called `College_Admissions`, either as one long file or as a series of modules in separate files. Using local or nested functions provides two other ways of organizing this program, retaining the modular organization of the top-down programming approach but also putting all the code into a single file rather than into multiple files. The nested-function approach is shown in Code 8.5.6. Because the functions are nested, they have access to the variables of the main program, so no arguments need to be passed to the functions or returned from them. The output is omitted because it is identical to that of Output 8.1.1.

Code 8.5.6:

```
% College_Admissions_Nested
function main;
clear all
clc
commandwindow

% Set constants
applications = 30;
max_admits_allowed = 10;
```

```

IQmean = 110;
IQsd = 20;
SATQmean = 500;
SATQsd = 100;
SATVmean = 500;
SATVsd = 100;
ECsd = 10;
GPAmean = 2.0;
GPAsd = 10;
% Use deal to initialize all these variables in one command
[IQ SATQ SATV GPA Acad EC Dist] = deal(zeros(applications,1));
Scores = [];
ScoresAndAcceptances = [];
criterion = [];

% Program sequence
Generate_Dummy_Scores;
Normalize_Scores;
Create_Scores_Matrix
Select_Students
Display_Results
return

function Generate_Dummy_Scores;
    IQ = IQmean + (randn(1,applications)) * IQsd;
    SATQ = SATQmean + (randn(1,applications)) * SATQsd;
    SATV = SATVmean + (randn(1,applications)) * SATVsd;
    GPA = GPAmean + (randn(1,applications)) * GPAsd;
    EC = abs(randn(1,applications)) * ECsd;
    Dist = abs(randn(1,applications));
    Acad = SATQ + SATV + 100 * GPA;
    return
end % Generate_Dummy_Scores

function Normalize_Scores;
    Acad = (Acad - min(Acad)) ./ (max(Acad)-min(Acad));
    EC = (EC -min(EC)) ./ (max(EC)-min(EC));
    Dist = (Dist - min(Dist)) ./ (max(Dist)-min(Dist));
    return
end % function Normalize_Scores

function Create_Scores_Matrix;
    % Create a Scores matrix, including, in the final column,
    % each student's total score
    Scores = [[1:applications]' Acad EC Dist];
    Scores(:,5) = Acad + EC + Dist;

```

```
        return
    end % Create_Scores_Matrix

    function Select_Students;
        % Admit the top max_admits_allowed students (plus
        % any ties)
        SortedScores = sortrows(Scores,-5);
        criterion = SortedScores(max_admits_allowed,5);
        SortedScores( : ,6) = 0;
        SortedScores((SortedScores(:,5) >= criterion),6) = 1;
        ScoresAndAcceptances = sortrows(SortedScores,1);
        return
    end % Select_Students;

    function Display_Results;
        % Display the results
        fprintf('App.\tAcad.\tExtra.\tDist.\tTotal\tAccept\n\n')

        fprintf('%4d\t%6.2f\t%6.2f\t%6.2f\t%6.2f\t%4d\n',...
            ScoresAndAcceptances)
        fprintf('\r')
        Students_Accepted = find(ScoresAndAcceptances(:,6));
        fprintf('Accepted Students:\n');
        fprintf('%3d',Students_Accepted);
        fprintf('\n')
        fprintf('Cutoff score: %5.03f\n', criterion);
        return
    end % function Display_Results

end % Function Main
```

The important take-home lesson from this set of nested functions is that every function has direct access to every variable in the main program. On the other hand, the initialization of variables cannot be delegated to a nested function because every variable referred to by one of the nested functions has to be referenced at some point in the main function. Any variable used *only* in a nested function would be invisible to the main and other nested functions.

A hierarchically organized program like this can serve as a good way of attacking a problem, for it can help you first focus on the major organization of operations that needs to be performed, and then let you focus your full attention on the individual operations within the subfunctions, one at a time.

There is one other point that is particularly relevant to behavioral scientists who depend on the consistency of timing in their programs for stimulus display or response detection. It is that the *first* time a function in an external file is called, it must be loaded into computer memory, which takes a fraction of a second for disk access and compilation. Once loaded, however, the function stays resident in memory, so there is no subsequent delay. For maximal consistency of timing,

then, it may be desirable to use local or nested functions in preference to external functions. In addition, it may be wise not to rely on the first trial of an experimental session with the expectation that it affords accurate timing (practice or warm-up effects for the participant aside).

8.6 Calling Functions Properly

It is worth taking a moment to emphasize the importance of calling functions properly. Not calling functions in the way they are designed, or mishandling the returned values, can lead to unexpected results.

Code 8.5.2 had the line, `[a b] = mean_and_trimmed_mean(x)`. It was essential to declare the *pair* of output values to be returned by the function because this particular function returned *two* values. If the call to `mean_and_trimmed_mean` did not list any output values or listed just one output value, only one value would be returned, namely, the first value returned by the function. The following code demonstrates what happens when different numbers of elements are indicated in calls to `mean_and_trimmed_mean`.

Code 8.6.1:

```
b = mean_and_trimmed_mean(x)
c = mean_and_trimmed_mean(x)
[d e] = mean_and_trimmed_mean(x)
[f g h] = mean_and_trimmed_mean(x)
```

Output 8.6.1:

```
b =
    3
c =
    3
d =
    3
e =
    2.5009
??? Error using ==> mean_and_trimmed_mean
Too many output arguments.
```

As the output of `d` and `e` shows, a function can return more than one value. On the other hand, a function cannot return more values than it was designed to.

8.7 Exploiting Recursive Functions

Sometimes a function performing some computation can profit from calling the same function itself. A function that calls itself is a *recursive* function. Recursion is useful in problems that do not have a more direct analytical solution or that would require a large number of nested `for` loops. The same result could be obtained by writing several nested `for` loops, but if the number of loops had to be changed, many lines of code would

have to be edited. In a recursive function, if the number of loops must be changed, only one number needs to be edited. Recursion can be useful when the number of loops may vary or tends to be very large.

The following code describes the general strategy of such a program, using a local function. The main program specifies how “deep” to carry the analysis. The recursive function calls itself repeatedly, at lower and lower levels, until it reaches the lowest level. In this example, the loop depth is 5, and the recursive call to `Doalooop` can be found between the `else` and the end of the `if` statement in `Doalooop`. The first call to `Doalooop`, in line 5 of the code, gets things started.

Code 8.7.1:

```
% Example of recursive routine, equivalent of 5 nested FOR loops
function recursiveFunction
loopdepth = 5;
thislevel = 0
Doalooop(thislevel, loopdepth);
% end of recursiveFunction

% ===== Local function called repeatedly by self
function Doalooop(thislevel,loopdepth);
thislevel = thislevel + 1;
enteringlevel = thislevel;
fprintf('Entering Level %d\n',thislevel');
if thislevel == loopdepth
    % Do what needs to be done in the 'innermost' loop;
    % necessary parameters would have been set
    % at each level of depth. . .
    fprintf('  At the lowest level\n');
    return
else
    % Not deep enough yet.
    % Set whatever parameters need to be set at this level
    % then call self recursively at a deeper level)
    Doalooop(thislevel,loopdepth);
end
leavinglevel = thislevel;
fprintf('Leaving Level %d\n',thislevel');
% end of Doalooop
```

Output 8.7.1:

```
thislevel =
         0
Entering Level 1
Entering Level 2
```

```

Entering Level 3
Entering Level 4
Entering Level 5
    At the lowest level
Leaving Level 4
Leaving Level 3
Leaving Level 2
Leaving Level 1

```

To make recursive programming less abstract, here is a recursive routine that performs a multidimensional search by setting four joints in a hypothetical two-dimensional stick figure to a succession of values. Many details are left out, but the overview is that each of four joints (trunk, shoulder, elbow, and wrist) takes on a value of -1 , 0 , and 1 , successively, so an exhaustive search of all three values for each of the four joints (3^4 , or 81 combinations in all) can be explored. Here is what a loop-based version of the program would look like.

Code 8.7.2:

```

clc
fprintf(['Joint values          T  S  E  W\n'...
'          _ _ _ _ \n']);
for trunkvalue = [-1:1]
    for shouldervalue = [-1:1]
        for elbowvalue = [-1:1]
            for wristvalue = [-1:1]
                fprintf(['Processing values'...
                    '%3.0f%3.0f%3.0f%3.0f\n'],...
                    trunkvalue, shouldervalue,...
                    elbowvalue, wristvalue);
            end
        end
    end
end

```

Output 8.7.2:

```

Joint values          T  S  E  W
Processing values    -1 -1 -1 -1
Processing values    -1 -1 -1  0
Processing values    -1 -1 -1  1
Processing values    -1 -1  0 -1
Processing values    -1 -1  0  0
Processing values    -1 -1  0  1
. . . 69 lines of output omitted . . .
Processing values     1  1  0 -1

```

```
Processing values 1 1 0 0
Processing values 1 1 0 1
Processing values 1 1 1 -1
Processing values 1 1 1 0
Processing values 1 1 1 1
```

When such a search involves many joints, or many values of each joint, the recursive routine is valuable. It is relatively easy to vary the number of levels (joints) and the number of joint values explored at each level because a single change in the recursive routine affects all the loops.

Code 8.7.3:

```
% Code 8_7_3.m
function recursiveKinematics
clc
loopdepth = 5;
thislevel = 0;
jointvalues = zeros(1,4);
fprintf(['Joint values          T S E W\n'...
'          _ _ _ _\n']);
Doalooop(thislevel,loopdepth, ...
{ 'Trunk','Shoulder','Elbow','Wrist'}, jointvalues);
return

% ===== Recursive function

function Doalooop(thislevel,loopdepth,joints,jointvalues);
thislevel = thislevel + 1;
enteringlevel = thislevel;
if thislevel == loopdepth
    fprintf(['Processing values' ...
            '%3.0f%3.0f%3.0f%3.0fn'],jointvalues);
    return
else
    for i = -1:1
        jointvalues(thislevel) = i;
        Doalooop(thislevel,loopdepth,joints,jointvalues);
    end
end

return
```

The output of this version is not shown because it would be identical to Output 8.7.2.

Another application for recursion is one in which you don't know ahead of time how many loops are required. Imagine writing a function to look for files that have certain characteristics. If you wrote a folder-scanning function (you could call it `ScanMyFolder.m`) to look at

a folder that itself contains an unknown number of sub-folders (as well as sub-folders of *those* folders), you could use the techniques described in Section 7.4 to look at each file in turn by reading a folder's directory using `mydir = dir`. Whenever one of the entries in the current directory was a sub-folder rather than a regular file (as indicated by `mydir(filecount).isdir == 1`) you could call the folder-scanning function to recursively scan the *sub-folder* with the sub-folder's name (`mydir(filecount).name`) as an argument, and then return to the next higher level in the recursion once you processed the last entry in each sub-folder.

You may be able to think of other examples of computation that would profit from recursive evaluation. Often, parameter estimation problems in computational modeling or simulation that do not yield a formal mathematical solution can be approached through recursion. MATLAB provides tools for this. The built-in MATLAB function `fminsearch`, for example, implements a recursive search for parameters that may best fit a model, using the so-called Nelder-Mead simplex direct search algorithm (see Press, Teukolsky, Vetterling, & Flannery, 2007). More details about `fminsearch` can be found in MATLAB's documentation.

8.8 Drawing on Previously Defined Functions Versus Creating Your Own

A final remark about functions is that sometimes you may have to decide between exploiting previously defined functions versus creating your own functions from scratch. Each approach has advantages and disadvantages.

MATLAB comes with a large number of built-in functions that have been optimized and extensively tested. People in the MATLAB programming community also provide functions for free on the MathWorks support site (www.mathworks.com/matlabcentral/). It is useful to draw on these sources if creating your own function seems daunting or needlessly time-consuming. In addition, carefully studying the code that others have developed has great heuristic value.

On the other hand, using other people's functions can leave you at their mercy. You may be stuck with code that has a bug in it or is difficult to verify to your satisfaction, or it may not quite address the problem you need to address, in which case you might spend more time trying to find a function that does what you want than generating it yourself. Do not shrink from writing your own functions. It is instructive to do so. The depth of your own understanding will increase if you write functions of your own design. You don't have to do so from scratch, however. You can also edit existing functions to turn them into ones you want.

8.9 Practicing Modules and Functions

Try your hand at the following exercises, using only the methods introduced so far in this book or given in the problems themselves.

Problem 8.9.1:

Write a function to convert any specified value of a normally distributed 1000×1 random sample, `mysample`, to a *z* score. The *z* score of such a value is its signed number of sample standard deviations away from the sample mean. What arguments will the function need to call? Save your function in an `.m` file, and call it from another script or the Command

window, with one argument, `mysample`. Then make it a nested function, and call it from the main function in your `.m` file. Test the function with random samples of different sizes, means, and standard deviations.

Problem 8.9.2:

In Problems 5.92 and 5.93 you were asked to identify participants who had mean reaction times greater than 500 ms and proportions correct greater than .65. If you solved the problem and followed the instruction to use material presented up to that point only, or information given in the problems themselves, you did so without creating a function. Now, write a function that takes as input these three variables: (1) the name of the matrix containing reaction times and proportions correct; (2) the reaction time cutoff; and (3) the proportion-correct cutoff. The function should return the following:

(1) `Identified_Participants`; (2) `OK_Scores`; (3) `Mean_of_OK_reaction_times`; and (4) `Mean_Proportion_Correct`.

Problem 8.9.3:

It would be desirable to apply the function you created in the last problem to a larger data set than the one given in Problem 5.9.2. You needn't collect actual data for this purpose. Instead, you can generate model data via simulation. Generate model data that reflect the following constraints: (1) There are 1,000 trials; (2) the probability of a correct response on any given trial is .90; (3, 4) reaction times in correct trials are drawn from a normal distribution with $\mu = 700$ ms and $\text{std} = 20$ ms; (5, 6) reaction times in incorrect trials are drawn from a normal distribution with $\mu = 600$ ms and $\text{std} = 80$ ms; reaction times less than 0 ms are undefined. Generate the model data based on the above constraints with a single function that has five arguments corresponding to the values of constraints 1–6, so you can run the function with a different set of constraint values in the future.

Problem 8.9.4:

Write a function to compute the probability of getting exactly k successes in n tries given the constraints outlined below (such as getting exactly four heads in 10 flips of a fair coin). Quoting from an August 31, 2006, entry in Wikipedia (http://en.wikipedia.org/wiki/Binomial_distribution), “if the random variable X follows the binomial distribution with parameters n and p , we write $X \sim B(n, p)$. The probability of getting exactly k successes is given by the probability mass function:

$$f(k, n, p) = \binom{n}{k} p^k (1-p)^{n-k}$$

for $k = 0, 1, 2, \dots, n$, where

$$\binom{n}{k} = \frac{n!}{k!(n-k)!}$$

Recall that $n!$ is called “ n factorial” and is equal to $1 \times 2 \times 3 \times \dots \times (n-1) \times n$. Likewise, $k!$ is called “ k factorial” and is equal to $1 \times 2 \times 3 \times \dots \times (k-1) \times k$. You might wish to start by

exploring `help factorial`, or write your own program to generate the factorial of an argument using a recursive function. How would you verify the accuracy of your function? Can you model the process using randomly generated data?

Problem 8.9.5:

Pascal's triangle is an arrangement of numbers such that the numbers in each row are generated from the sum of the two numbers directly above and to the left. The first four lines of the triangle are:

```
1
1 1
1 2 1
1 3 3 1
```

Each line can be generated from the immediately prior line as: `nextrow = [thisrow[0] + [0 thisrow]`; thus the next line in the above list would be `[1 3 3 1 0] + [0 1 3 3 1]`, or `[1 4 6 4 1]`. Implement the generation of Pascal's triangle, starting with `thisrow = 1`, by a recursive function `nextrow = PascalOf(thisrow)`, where the criterion for stopping the recursion is `length(thisrow) >= 12`. After each iteration, print both each line and the sum of the values in that line. Is there any regularity in the growth of the sum from line to line?

Problem 8.9.6:

Find a problem that you solved for a prior chapter that would profit from being organized in modular fashion, and rewrite it using either local functions (in the same file as a main function) or nested functions. Look for programs where you have had to execute the same computation repeatedly, varying only one or two arguments of the computation.

Problem 8.9.7:

In Problem 7.7.4 you were asked to write a program to administer a computerized questionnaire on a topic of interest to you. You were asked to use a structure data type and to allow participants to answer with whole sentences or phrases for at least some items. You were asked to save the data in an external file, and you were asked to record the times taken to answer the questions. Make this program modular and, having done so, take advantage of that modularity to pursue different lines of questions depending on participants' answers to particular questions.

9. Plots

This chapter covers the following topics:

- 9.1 Deciding to plot data and, for starters, generating a sine function
- 9.2 Controlling axes
- 9.3 Controlling the appearance of plotted points and lines
- 9.4 Having more than one graph per plot and more types of points and lines
- 9.5 Getting and setting properties of plotted points
- 9.6 Adding xlabel, ylabel, and title
- 9.7 Adding legends
- 9.8 Adding text
- 9.9 Fitting curves
- 9.10 Creating and labeling subplots and turning grids, boxes, and axes on and off
- 9.11 Exploiting matrix assignments to merge subplots
- 9.12 Getting and setting properties of axes
- 9.13 Plotting data points with error bars
- 9.14 Generating polar and compass plots
- 9.15 Generating histograms
- 9.16 Generating bar graphs
- 9.17 Saving, exporting, and printing figures
- 9.18 Generating other kinds of graphs and getting and setting figure properties
- 9.19 Practicing plots

The commands that are introduced and the sections in which they are premiered are:

<code>close</code>	(9.1)
<code>clf</code>	(9.1)
<code>figure</code>	(9.1)
<code>plot</code>	(9.1)
<code>shg</code>	(9.1)
<code>sin</code>	(9.1)
<code>axis</code>	(9.2)
<code>xlim</code>	(9.2)
<code>ylim</code>	(9.2)
<code>'g- '</code>	(9.3)
<code>'bo '</code>	(9.3)
<code>cos</code>	(9.4)
<code>hold</code>	(9.4)

color	(9.5)
get	(9.5)
markeredgecolor	(9.5)
markerfacecolor	(9.5)
markersize	(9.5)
title	(9.6)
xlabel	(9.6)
ylabel	(9.6)
legend	(9.7)
text	(9.8)
polyfit	(9.9)
box	(9.10)
get(h)	(9.10)
grid	(9.10)
set(h, 'Position')	(9.10)
subplot	(9.10)
get(gca)	(9.12)
set	(9.12)
errorbar	(9.13)
compass	(9.14)
polar	(9.14)
brighten	(9.15)
colormap	(9.15)
hist	(9.15)
bar	(9.16)
barh	(9.16)
loose	(9.17)
print	(9.17)
saveas	(9.17)
feather	(9.18)
get(0, 'Screensize')	(9.18)
get(gcf)	(9.18)
pie	(9.18)
plotyy	(9.18)
quiver	(9.18)

<code>set(gcf, 'Position')</code>	(9.18)
<code>stairs</code>	(9.18)
<code>stem</code>	(9.18)

9.1 Deciding to Plot Data and, for Starters, Generating a Sine Function

As mentioned in the Preface, one of MATLAB's most attractive features is that it lets you easily generate data plots and other graphics. The fact that MATLAB offers many options for plotting accounts for the fact that this is one of the longest chapters in this book.

The first step in creating a data plot is deciding whether you actually need one. A well-designed data plot lets you see trends in your data or lets you (or perhaps *compels* you to) lower your expectations about such trends. If plotting the data shows that the data look more like a blizzard than a line, that fact may cause you to rethink a hypothesis that predicted a strong relationship.

Creating well-designed data plots takes practice, but MATLAB provides a convenient medium for honing your graphing skills


To create your first data plot, follow Code 9.1.1. Here we start with a clean slate by clearing all variables, using `clear all`. Then we close all currently active figures using `close all`. The currently active figure is where new plots and graphics will appear. Using the `close all` command is advisable unless you want to add a plot to an existing figure. To emphasize that last point and to say it another way, if you *want* to add a plot to an existing figure, *don't* close that figure. The syntax for closing a figure of your choice, and so for *not* closing a figure of your choice, is given in the next paragraph. To clear the currently active figure without closing it (i.e., keeping that figure window open but wiping it clean), you can use the `clf` command.

Suppose you want to create figure number 1. You can do so using the command `figure(1)`. MATLAB assumes that the first figure number is 1, so you could just as well have written `figure`. However, it's useful to know about figure numbers in general in case you want to generate series of figures, such as `figure(2)`, `figure(3)`, and so on. This can make it easier for you to refer back to the figures later in your program or to find the saved versions easily in your disk directory. We will explain how you can save figures later in this chapter. By having individually numbered figure windows, you can close those windows selectively. For example, if you want to close figure (3), you can say `close (figure(3))` or `close(3)` for short. If the current value of `f` is 3, you can say `close(f)` instead.

Code 9.1.1 lets you generate the graph of a `sin` (pronounced “sine”) function. Here is the code and the output it produces, shown via the `shg` command.

Code 9.1.1:

```
% Code_9_1_1.m
clear all
close all
figure(1)
theta_rad = linspace(0,4*(2*pi),100);
plot(theta_rad,sin(theta_rad));
shg
```

Output 9.1.1:

Let's decode this program. First, it's useful to recall that the sine function is related to the vertical position of the end of a unit radius of a circle as it rotates (counterclockwise, by convention) starting from 0 degrees (toward the right, again by convention). MATLAB assumes that angles increase in the counterclockwise angular direction, as just implied, and that an angle whose value is 0 is associated with the straight line extending from the center of the circle to the right, again as just implied. You know that there are 360 degrees in a circle, but there are also 2π "radiuses" or radians. In other words, if you take a string whose length is the same as the radius of the circle and ask how many of those strings fit exactly around the rim of the circle, the answer is 2π . The value π is just the ratio of the circumference of a circle to its diameter. Because the radius of a circle is half the length of its diameter, the ratio of the circumference of a circle to its radius is 2π . Sine is 1.0 when the radius points straight up (i.e., when it has rotated 90 degrees or $\pi/2$ radians). Sine is -1.0 when the radius points straight down i.e., when it has rotated 270 degrees or $3\pi/2$ radians).

MATLAB uses radians rather than degrees in almost all of its trigonometric calculations. Pay close attention to that last statement! Forgetting it, thinking that angles are measured in degrees, can cause you a lot of grief. For this reason, a book we admire for its wisdom about the naming of variables (Johnson, 2011) recommends that when you use units of

measurement, append a suffix (`'_rad'` or `'_deg'`) to the variable name. That way, a statement like `theta_rad = theta_deg * pi/180` will be clear. By the same token, if your program uses both English and metric units of length, you might use a variable name `width_cm` or `width_in`, depending on the units used. In case you think this is a small point, just a matter of concern for neophyte programmers using a book like this to learn or be reminded of the basics of programming, consider the sad fate of the Mars Climate Orbiter, which failed, despite millions of dollars going into its development, because of a mix-up concerning the unit of measure of its optical measurement device (http://en.wikipedia.org/wiki/Mars_Climate_Orbiter).

For the graph generated here, the radius is rotated four times, taking 100 equal steps along the way. To plot the function, we define a matrix `theta_rad` as an array of 100 elements, each representing an angle, linearly spaced between a minimum value of 0 and a maximum value of $4 \times 2 \times \pi$. The `sin` function evaluated from 0 to 8π represents four complete turns of the radius, or four cycles of the sine wave.

To plot `sin(theta_rad)` as a function of `theta_rad`, we use the `plot` command. Keep in mind that the first argument provided to `plot` is the array for the horizontal axis, or *abscissa*, of the graph. The second argument is the array for the vertical axis, or *ordinate*, of the graph. (“Abscissa” is a more general term than “x-axis,” and “ordinate” is a more general term than “y-axis.”)

As you can see in Output 9.1.1, the function `sin(theta_rad)` oscillates around 0 with a maximum of 1 and a minimum of -1. This is because `sin(theta_rad)` is obtained by taking the height (the vertical position) of the end of the radius after a given rotation `theta_rad` and dividing that height (which varies with the rotation) by the length of the radius, which is fixed. When $\theta = 0$ radians (also 0 degrees), the height of the end of the radius is 0 times the radius; hence $\sin(\theta) = 0$. (Note that θ is the Greek letter for theta, the standard mathematical notation for an angle.) When $\theta = 1/4 \times 2 \times \pi = \pi/2$ radians (or 90 degrees), the height of the end of the radius is equal to +1 times the radius; hence $\sin(\pi/2) = 1$. When $\theta = 2/4 \times 2 \times \pi = \pi$ radians (or 180 degrees), the height of the end of the radius is again 0 times the radius; hence $\sin(\pi) = 0$. When $\theta = 3/4 \times 2 \times \pi = 3\pi/2$ radians (or 270 degrees), the height of the end of the radius is -1 times the radius; hence $\sin(3\pi/2) = -1$. Finally, when $\theta = 4/4 \times 2 \times \pi = 2\pi$ radians (or 360 degrees), the height of the end of the radius is once again 0 times the radius; hence $\sin(2\pi) = 0$. The sine function can keep on going forever, ascending and descending in a perfectly periodic fashion, which is why the sine is a so-called periodic function.

9.2 Controlling Axes


Output 9.1.1 isn't as pretty as it might be. One problem is that the curve ends abruptly, leaving a lot of room to spare. It would be nice to fix this. You can do so by defining the range of the axes, using `axis`. The `axis` function requires four values: the smallest and largest values on the horizontal axis, and the smallest and largest value on the vertical axis. In the code that follows, these four values are defined generically using the built-in functions `min` and `max` (see Chapter 3). Because the four elements constitute

a matrix, they must be enclosed in brackets, as in any standard MATLAB matrix with more than one element.

Code 9.2.1:

```
% code_9_2_1.m
figure(2)
y = sin(theta_rad);
plot(theta_rad,y);
axis([min(theta_rad) max(theta_rad) min(y) max(y)]);
shg
```


Output 9.2.1:



This graph looks better than its predecessor. However, it could be even prettier if we defined the axes so there were some “space to breathe” above and below the plotted points. The next program generalizes the preceding code by adding more information concerning the minima and maxima for the x and y axes. It also illustrates another way of specifying those values that does not require the use of the `axis` command, not that there is anything wrong with that command. The alternative method is to use `xlim` and `ylim`. These functions have the advantage that they can be used independently of one another, allowing you to specify the limits of the x axis only or the y axis only. The `axis` command, by contrast, forces you to specify the limits of x and y .

Code 9.2.2:

```
% code_9_2_2.m
figure(3)
x = theta_rad;
plot(x,y);
x_offset = 1;
y_offset = .2;
xlim([min(x)-x_offset, max(x+x_offset)]);
ylim([min(y)-y_offset, max(y+y_offset)]);
shg
```


Output 9.2.2:


9.3 Controlling the Appearance of Plotted Points and Lines

We can control the way plotted points appear. Adding 'g-' to the `plot` command tells MATLAB to connect the points with a green (g) line (-). Because we are just adding information to the figure, there is no need to specify `xlim` and `ylim` again, just as there is no need to specify `x` and `y` again because these values are active, owing to the fact that they haven't been cleared, nor have we quit or restarted MATLAB, which would have cleared all active variables and figures.

In the program below, we use another new command, `hold on`, which tells MATLAB to maintain the already plotted figure when new material is added to it. In this case, blue o's are added to the graph. Note that these are blue letter-o's, not blue zeros. To see the o's in color rather than the grayscale used in this book, go to the book's website (www.routledge.com/9780415535946), or run the program by typing it into your Command window.

Code 9.3.1:

```
% code_9_3_1
figure(4);
plot(x,y,'g-');
hold on;
plot(x,y,'bo');
shg;
```

Output 9.3.1:

9.4 Having More Than One Graph per Plot and More Types of Points and Lines


The `hold` on command is especially useful when you want to have more than one graph per plot. The program below shows how you can achieve this. First tell MATLAB to create a new figure, `figure(5)`. Then issue a command to plot `y` against `theta_rad` using green `o`'s and green line segments. Notice that the color and shape of the points as well as the line segments are indicated in a single command, `plot(theta_rad, y, 'go-')`.

To see what can be achieved with the `hold` on command, add a second curve to `figure(5)`. Besides plotting `sin(theta_rad)` as a function of `theta_rad`, also plot `cos(theta_rad)` as a function of `theta_rad`. The command `cos` is a built-in MATLAB function, as is `sin`. The function `cos(theta_rad)` takes the horizontal position of the end of the radius at a given angle `theta_rad` and divides that horizontal position by the length of the radius. As seen below, `cos(theta_rad)` is plotted as a function of `theta_rad` using blue line segments and blue squares (`'b-s'`).

Code 9.4.1:

```
% code_9_4_1.m
figure(5)
theta_rad = 0:.1:2*pi;
y = sin(theta_rad);
plot(theta_rad, y, 'go-'); hold on;
y = cos(theta_rad);
plot(theta_rad, y, 'b-s');
```

Output 9.4.1:



Code 9.4.2 shows that MATLAB plots can include a variety of colors and shapes of points and lines. So far you have used blue and green circles and squares, as well as blue and green line segments. By typing `help plot`, you can learn about the full range of plotting options that MATLAB provides. Output 9.4.2 is an excerpt from the information that is returned when you type `help plot`.

Code 9.4.2:

```
| help plot
```

Output 9.4.2:

Various line types, plot symbols and colors may be obtained with `PLOT(X,Y,S)` where `S` is a character string made from one element from any or all the following 3 columns:

b	blue	.	point	-	solid
g	green	o	circle	:	dotted
r	red	x	x-mark	-.	dashdot
c	cyan	+	plus	--	dashed
m	magenta	*	star	(none)	no line
y	yellow	s	square		
k	black	d	diamond		
		v	triangle (down)		
		^	triangle (up)		
		<	triangle (left)		
		>	triangle (right)		
		p	pentagram		
		h	hexagram		


For example, `PLOT(X,Y,'c+:')` plots a cyan dotted line with a plus at each data point; `PLOT(X,Y,'bd')` plots blue diamond at each data point but does not draw any line.

By relying on the foregoing information, you can specify other colors, shapes, and line types. The following program illustrates this fact and also reveals another useful feature of plotting, namely, that it is possible to tell MATLAB to generate two (or more) graphs with one `plot` command. Here, in one `plot` statement, you indicate that you want to plot $\sin(x)$ against x using cyan plus signs connected by a dotted line, and also that you want to plot $\cos(x)$ against x using red diamonds not connected by a line. Both instructions can be given in one line of code. There is no particular advantage to writing the code this way, except to consolidate it, so it is simply a matter of personal preference.

Code 9.4.3:

```
figure(6)
theta = 0:.1:2*pi;
plot(theta,sin(theta),'c+:',theta,cos(theta),'rd');
shg
```

Output 9.4.3:




9.5 Getting and Setting Properties of Plotted Points

You can control the size of plotted points using one of their properties: `markersize`. Setting `markersize` to 12 yields larger circles than in the previous outputs.

Code 9.5.1:

```
figure(7)
x = 0:.1:2*pi;
plot(x,sin(x),'ro-', 'markersize',12);
xlim([min(x)-x_offset, max(x+x_offset)]);
ylim([min(y)-y_offset, max(y+y_offset)]);
box on
shg
```

Output 9.5.1:


How do you find out about properties such as `markersize`? You can do so with one of the most useful commands in MATLAB: `get`.

Code 9.5.2 shows how you can `get` the properties of a plot similar to the one above. The third line of Code 9.5.2 shows how the `get` command is used. `get` is a function whose argument (in this case, `h`) is a set of parameters associated with the `plot` function, called in the second line of Code 9.5.2.

Output 9.5.2 includes text returned via `get(h)`. The graph reveals two things—first, that $\sin(x)$ plotted as a function of $\cos(x)$ yields a circle, and second, that the actual size of plotted points depends on the type of point as well as the value of `markersize`. Compare the size of the points in Output 9.5.2 with the size of the points in Output 9.5.1, where the value of `markersize` is the same but the types of plotted points are different.

Code 9.5.2:

```
figure(8)
x = 0:1:2*pi;
h = plot(cos(x), sin(x), 'r.', 'markersize', 12);
axis equal
get(h)
```

Output 9.5.2:

```

h =
    DisplayName: ''
    Annotation: [1x1 hg.Annotation]
    Color: [1 0 0]
    LineStyle: 'none'
    LineWidth: 0.5000
    Marker: '.'
    MarkerSize: 12
    MarkerEdgeColor: 'auto'
    MarkerFaceColor: 'none'
    XData: [1x63 double]
    YData: [1x63 double]
    ZData: [1x0 double]
    BeingDeleted: 'off'
    ButtonDownFcn: []
    Children: [0x1 double]
    Clipping: 'on'
    CreateFcn: []
    DeleteFcn: []
    BusyAction: 'queue'
    HandleVisibility: 'on'
    HitTest: 'on'
    Interruptible: 'on'
    Selected: 'off'
    SelectionHighlight: 'on'
    Tag: ''
    Type: 'line'
    UIContextMenu: []
    UserData: []
    Visible: 'on'
    Parent: 173.0519
    XDataMode: 'manual'
    XDataSource: ''
    YDataSource: ''
    ZDataSource: ''

```

By knowing the properties of a plotted figure, you can set the properties you want. For example, you can control the `markerfacecolor` and `markeredgecolor` of plotted points, as shown below. The colors that appear are likely to be more vivid on the screen or website than on this printed page.

Code 9.5.3:

```

figure(9)
x = theta_rad;
plot(x,y,'g-');
hold on


```

```

x_offset = 0;
y_offset = .2;
axis([min(x)-x_offset, max(x)+x_offset, ...
      min(y)-y_offset, max(y+y_offset)]);
plot(x,y,'o', 'color','r','markersize',6,...
      'markeredgecolor','k','markerfacecolor','r');

```

Output 9.5.3:



Changing the `markerfacecolor` and the `markeredgecolor` of plotted points is just one thing that can be done by varying figure properties. The method illustrated in Code 9.5.3 can be generalized to other properties of interest. For example, `color` can be specified as shown in Code 9.5.3, where `'color'` is followed by a single letter code such as `'r'`. Alternatively `'color'` can be followed by a 1×3 matrix, such as `[1 0 0]`. The first number is the value of red, the second number is the value of green, and the third number is the value is blue. It is easy to remember this order by memorizing the letters RGB. A further mnemonic is to think of RGB as the initials of the fictional character Roy G. Biv, or make up your own personally meaningful mnemonic. Setting each of the three numbers associated with `'color'` to values between 0 and 1 will let you create almost any color you want. Only values between 0 and 1 are permissible as values for `'color'` because each is a proportion of the maximum for that color, and proportions can only range from 0 to 1.

9.6 Adding Xlabels, Ylabels, and Titles

You can generate a graph like the one shown in Output 9.6.1 by adding an `xlabel`, a `ylabel`, and a `title`.

Code 9.6.1:

```

figure(10)
plot(x,y,'g-');


```

```

hold on
x_offset = 0;
y_offset = .2;
axis([min(x)-x_offset, max(x)+x_offset, ...
      min(y)-y_offset, max(y)+y_offset]);
plot(x,y,'o','color','r','markersize',6,...
      'markeredgecolor','k','markerfacecolor','r');
xlabel('Time');
ylabel('Happiness');
title('Life has its ups and downs');

```

Output 9.6.1:



9.7 Adding Legends

You can also add a legend to a graph, as in the following example, where hypothetical learning curves are generated for subjects in four conditions, *c1*, *c2*, *c3*, and *c4*, who try to recall the same items after the items are presented in identical fashion in successive trials. The learning curves are based on the idea that the four conditions have different asymptotes and that the rate at which the asymptotes are approached diminish the longer the experiment continues.

In creating a legend, you assign strings to each curve. The order of the strings should correspond to the order in which the data are plotted. This is why the order of plotting the curves below is “backwards.” In the code below, the curves are plotted in an order that ensures good stimulus-response compatibility between the items in the legend and the curves themselves (the higher the legend, the higher the curve). The arguments at the end of the `legend` command tell MATLAB where the legend should be placed. Other options for legend placements are available. For more information about this, type `help legend` at the MATLAB command line.

Code 9.7.1:

```

figure(11)
max_learn = [10 11 12 13];
trial = [1:10];

```




```

c1 = max_learn(1) - exp(-trial);
c2 = max_learn(2) - exp(-trial);
c3 = max_learn(3) - exp(-trial);
c4 = max_learn(4) - exp(-trial);

hold on
plot(trial,c4,'g-^');
plot(trial,c3,'m--<');
plot(trial,c2,'b-.>');
plot(trial,c1,'k:v');
legend('Group 4','Group 3',...
      'Group 2','Group 1',...
      'Location','EastOutside');

```

Output 9.7.1:

9.8 Adding Text

You can add text to a figure, as will be shown in the following examples, where a power function and an exponential function are plotted. A power function is one in which the independent variable is raised to some numerical power. The time it takes to perform a task is sometimes said to diminish with practice in a way that follows a power function. An exponential function is one in which the independent variable is itself part of the exponent to which some quantity is raised, as in the learning example above.

In the code that follows, we label the two curves using the `text` command. Note that the `text` command has three arguments: (1) the horizontal position where the text begins; (2) the vertical position where the text begins; and (3) the actual text string. In Code 9.8.1, we add a vertical offset and a horizontal offset to avoid crowding the text onto the curves. The vertical offset and the horizontal offset were found through trial and error. Note that the units governing the placement of the text are in the units of the particular axis we are using. Text drawn at (20, .9) would appear at the top right of the graph, for example.

Code 9.8.1:

```

figure(12)
a = 1;                % starting value
b = .5;               % rate parameter
xx = [0:20];
vert_offset = .05;
hor_offset = .50;


y_power = a * xx.^-b;
y_exp = a * exp(b*-xx);

hold on
box on
plot(y_power, 'mo-');
plot(y_exp, 'kd-');

hor_p = xx(5) + hor_offset;
vert_p = y_power(5) + vert_offset;
text(hor_p,vert_p, 'Power function');

hor_e = xx(6) + hor_offset;
vert_e = y_exp(6) + vert_offset;
text(hor_e,vert_e, 'Exponential function');

```

Output 9.8.1:

Sometimes, you may be interested in displaying *only* text. The next example implements the Stroop test (see MacLeod, 1991, for a review) to demonstrate interference in color-naming performance when word and font color are incompatible because they specify different responses. Potential words are first enumerated in the cell array `words`, and potential colors are listed in the variable `colors`. Note that a constant like `'rgbk'` can be treated as either a string or as an array of characters. In this instance, the string property is exploited to define the colors, and the array property is used to select the single character that represents the color of the word to be displayed in a particular trial. Each test word and trial type are displayed on the screen for 2 seconds to give the participant time to report the color of the letters. In the example, the word *Green* is in the color *black*, so it's an incompatible trial: the subject must ignore “Green” and say “black.”

Every graphics object (figure, line, text, data, etc.) can be assigned a variable called a *handle*. A handle can be used to manipulate the graphic object in a number of ways, as seen below. Here, we use the handles of the text items to remove them from the screen, by using `delete` to remove each object, (`delete([myhandle otherhandle])`). In general, deleting the handle of a graphic object in a figure removes it from the figure. Said in another, more vivid way, you can grab any object by its handle and toss it.

Code 9.8.2:

```
clear
close all;
figure('Name','Stroop Test')
words = {
    'Red'
    'Green'
    'Blue'
    'Black'
};
colors = ['rgbk'];
shg
text(.1,.8,sprintf(...
    ['Report the COLOR of the text\n', ...
    'as quickly as you can!']), ...
    'FontSize',18)
axis off
for t = 1:20
    w = randi(4);
    c = randi(4);
    myWordHandle = text(.2,.5,...
        char(words(w)), 'Color', colors(c), ...
        'FontSize',48);
    if w == c
        conditionstring = sprintf('Compatible trial');
    else
        conditionstring = sprintf('Incompatible trial');
    end
    myConditionHandle = text(.2,.2,conditionstring);
    pause(2)
    delete([myWordHandle myConditionHandle])
    pause(1)
end
```

Output 9.8.2:

Report the COLOR of the text
as quickly as you can!

Green

Incompatible trial

9.9 Fitting Curves

Behavioral scientists often fit curves to observed points. One way to do this is to use the `polyfit` function. This function lets you fit a polynomial function to data. A polynomial function of a variable x is a sum of terms consisting of a coefficient, often called a_0 , times x raised to the 0 power, plus another coefficient, often called a_1 , times x raised to the 1 power, plus another coefficient, often called a_2 , times x raised to the 2 power, all the way up to a coefficient, often called a_n , times x raised to the n power:

$$y = (a_0 * x^0) + (a_1 * x^1) + (a_2 * x^2) + \dots + (a_n * x^n)$$

Because any value raised to the 0 power is 1, $x^0 = 1$, in which case $(a_0 * x^0) = a_0$. Note that n defines the “order” of the polynomial.

In the following example, a set of dummy data is created based on a new matrix x , which runs from -20 to $+20$. To create the dummy data, we put each value of x through a second-order polynomial function to yield y , and then we add normally distributed random numbers to y , scaled by a coefficient arbitrarily called `randn_coeff`.

The first time we fit a curve to these data, we find a matrix of coefficients, called `fitted_coefficients`, which allows for a best fit of a first-order polynomial function. This is done with `polyfit(x, y, 1)`. The last term, 1, defines the order of the polynomial. A polynomial of order 1, or a “first-order polynomial,” is also called a *linear* equation. The best-fitting coefficients in this example are used to generate a matrix of theoretical values called `y_hat1`.

Code 9.9.1:

```
clear x y
a3 = 0;
a2 = 1;
a1 = 1;
a0 = 0;
```

```

x = [-20:20];
randn_coeff = 60;


y = a3*x.^3 + a2*x.^2 + a1*x.^1 + a0*x.^0;
r = rand(length(y))*randn_coeff;
r = r(1,:);
y = y + r;

fitted_coefficients = polyfit(x,y,1);
y_hat1 = fitted_coefficients(1)*x.^1 + ...
          fitted_coefficients(2)*x.^0; %apply polyfit
          % coefficients to x

figure (13)
hold on
plot(y, 'bo');           % show original data
plot(y_hat1, 'r-');      % show fitted points joined by a line
xlim([0 length(x)]);
box on                   % put a box around the graph
c = corrcoef(y, y_hat1);
message = ['Straight line fit: r^2 = ', num2str(c(1,2)^2,3)];
title(message);

```

Output 9.9.1:



As you can see, the fit isn't very good. The proportion of variance, r^2 , accounted for by the linear function is only .0046. To find r^2 (also known as the *coefficient of determination*), we computed the correlation matrix, arbitrarily called *c*, between *y* and *y_hat1* using `corrcoef`. Then we squared the element in the first row and second column of *c* to obtain r^2 (or we could have equally well squared the element in the second row and first column of *c*). To convert the value of r^2 to a string, suitable for presentation with the `title` command, we used the `num2str` command. The final term in the `num2str` command defined the number of significant figures.

Next, we seek a better fit with a second-order polynomial, also called a *quadratic* equation. We find a matrix of coefficients, arbitrarily called *pp2*, that allows for a best fit of a


second-order polynomial function. This is done using the command `polyfit(x,y,2)`. The coefficients are used to generate a matrix of theoretical values called `y_hat2`.

Code 9.9.2:

```
fitted_coefficients = polyfit(x,y,2);
y_hat2 = fitted_coefficients(1)*x.^2 + ...
         fitted_coefficients(2)*x.^1 + ...
         fitted_coefficients(3)*x.^0;

figure (14)
hold on
plot(y,'bo');           % show original data
plot(y_hat2,'r-');      % show fitted points joined by a line
xlim([0 length(x)]);
box on                  % put a box around the graph
c = corrcoef(y, y_hat2);
message = ['Quadratic fit: r^2 = ',num2str(c(1,2)^2,3)];
title(message);
```

Output 9.9.2:



The quadratic equation provides a much better fit to the data. The proportion of variance accounted for by the quadratic function exceeds .98.

9.10 Creating and Labeling Subplots and Turning Grids, Boxes, and Axes On and Off

You can generate several subplots within a figure using MATLAB's `subplot` function. This function has three arguments. The first is the number of subplot rows. The second is the number of subplot columns. The third is the number of the subplot that is about to be plotted, where the number increases from left to right and from top to bottom. The subplot command specifies where in the figure any new axes will be generated.

In the example that follows, we generate a 4×1 matrix of subplots. The first subplot, designated by `subplot(4,1,1)`, has the property that a grid is on. The second subplot, designated by `subplot(4,1,2)`, has the property that a box surrounds the graph. The third subplot, designated by `subplot(4,1,3)`, has the property that there is no axis. The fourth subplot, designated by `subplot(4,1,4)`, forces the graph to be square.

Note that `subplot` does not actually plot data. The `plot` command does this and is issued after the `subplot` command informs MATLAB which particular subplot is to be plotted next. Suffice it to say that the commands shown here for plotting the same data in different ways work even when subplots are not being used or, said another way, when the implicit subplot command is `subplot(1,1,1)`.

Code 9.10.1:

```
figure(15)
x = linspace(0,8*pi,100);


subplot(4,1,1)
plot(cos(x), 'r.', 'markersize', 12);
grid on

subplot(4,1,2)
plot(cos(x), 'r.', 'markersize', 12);
box on

subplot(4,1,3)
plot(cos(x), 'r.', 'markersize', 12);
axis off

subplot(4,1,4)
plot(cos(x), 'r.', 'markersize', 12);
axis square
```

Output 9.10.1:




When you use subplots, you may wish to label them efficiently. The next example shows how you can use `xlabel` to label the abscissa of two graphs. Code 9.10.2 creates a graph with two subplots.

Code 9.10.2:

```
figure(15)
clf;
subplot(1,2,1)
plot(cos(x),'r.','markersize',12);
grid on

subplot(1,2,2 )
plot(cos(x),'r.','markersize',12);
grid on
labelhandle = xlabel('Time(secs)')
```

Output 9.10.2:



Recall from Section 9.8 that every graphic object can have a *handle*, by which it can be manipulated in numerous ways. Code 9.10.3 adjusts the `xlabel`, whose handle is `labelhandle`, moving it to the left by 65 units and up by .05 units so it now applies to both subplots. The position is specified relative to a particular subplot, so the label “belongs” to the subplot on the right. The `labelposition` variable obtained using `get(labelhandle, 'position')` has three values, which can be useful for making a three-dimensional graph (see Section 10.8). You can ignore the third value if you are plotting in two dimensions, as in all the cases described here so far.


Code 9.10.3:

```
labelposition = get(labelhandle, 'Position')
labelposition(1) = labelposition(1) - 65;
```



```
labelposition(2) = labelposition(2) + .05;
labelposition
set(labelhandle, 'Position', labelposition, 'Fontsize', 18);
```

Output 9.10.3a:



Output 9.10.3b:

```
labelposition =
    49.7312   -1.1316    1.0001
labelposition =
   -15.2688   -1.0816    1.0001
```

9.11 Exploiting Matrix Assignments to Merge Subplots

You can merge subplots to enjoy considerable flexibility in the way your subplots appear. You can do this by building on methods covered in Chapter 3 for addressing different elements of a matrix. In the case of a matrix, we used the $r \times c$ rule (rows, then columns). You can address subplots the same way.

In the example that follows, a figure is created with a large title across the top, occupying subplots 1 and 2 of the 4×2 matrix of subplots to be drawn. Among the other subplots to be drawn, you can generate a graph in matrix positions 5 and 7.

Before showing the code used to generate the subplots (Code 9.11.1), it is worth mentioning that some of the features of the code were based on trial and error. For example, the number of spaces before the word `Banner` was adjusted by trying out different numbers of spaces, and the value of `.90` in the call to `text_in_box` for panel C was changed from the value of `.80` used in all the other panels because `.8` didn't result in as nice an appearance as we wanted. Trial

and error adjustment of parameters is often the most expedient, if not the most elegant, method of parameter specification. In what follows, we present Code 9.11.1 followed immediately by Code 9.11.2 before Output 9.11.1, just because Code 9.11.2 is a function used by Code 9.11.1.

Code 9.11.1:


```
function main
figure(16);
clf
clear x y
x = [1:10];
y = x + 1;
subplot(4,2,1:2); % In the 4 rows and 2 columns of subplots,
                  % subplots 1 and 2

xlim([0 1]);
ylim([0 1]);
axis off
text(-.05,.05,'                                A Banner Year',...
     'fontsize',24);
subplot(4,2,3); % In the 4 rows and 2 columns of subplots,
                % subplot 3
plot(x,y,'k')
text_in_box(.05,.80,'A')
subplot(4,2,4); % In the 4 rows and 2 columns of subplots,
                % subplot 4
plot(x,y,'k')
text_in_box(.05,.80,'B')
subplot(4,2,[5 7]); % In the 4 rows and 2 columns of subplots,
                   % subplots 5 and 7
plot(x,y,'k')
text_in_box(.05,.90,'C')
subplot(4,2,6); % In the 4 rows and 2 columns of subplots,
               % subplot 6
plot(x,y,'k')
text_in_box(.05,.80,'D')
subplot(4,2,8); % In the 4 rows and 2 columns of subplots,
               % subplot 8
plot(x,y,'k')
text_in_box(.05,.80,'E')
```

Code 9.11.2:

```
function text_in_box(x_place,y_place,s)

xs = xlim;
ys = ylim;
text(x_place*xs(2),y_place*ys(2),s);
```

Output 9.11.1:**9.12 Getting and Setting Properties of Axes**

Much as you can get the properties of plotted points by using the `get` function, you can get the properties of the axes of graphs with the `get (gca)` command. `gca` denotes the properties (or handles) of the current axes. Here, `get (gca)` is issued after running Code 9.11.1. The last axis plotted (and so, the current axis) was panel E. Here are all the attributes of panel E.

Code 9.12.1:

```
| get (gca)
```

Output 9.12.1:

```
ActivePositionProperty = position
ALim = [0 1]
ALimMode = auto
AmbientLightColor = [1 1 1]
Box = on
CameraPosition = [5 7.5 17.3205]
CameraPositionMode = auto
CameraTarget = [5 7.5 0]
CameraTargetMode = auto
CameraUpVector = [0 1 0]
CameraUpVectorMode = auto
CameraViewAngle = [6.60861]
CameraViewAngleMode = auto
CLim = [0 1]
CLimMode = auto
Color = [1 1 1]
CurrentPoint = [ (2 by 3) double array]
ColorOrder = [ (7 by 3) double array]
```

```

DataAspectRatio = [5 7.5 1]
DataAspectRatioMode = auto
DrawMode = normal
FontAngle = normal
FontName = Helvetica
FontSize = [10]
FontUnits = points
FontWeight = normal
GridLineStyle = :
Layer = bottom
LineStyleOrder = -
LineWidth = [0.5]
MinorGridLineStyle = :
NextPlot = replace
OuterPosition = [0.534263 0.0790476 0.409654 0.20298]
PlotBoxAspectRatio = [1 1 1]
PlotBoxAspectRatioMode = auto
Projection = orthographic
Position = [0.570341 0.11 0.334659 0.157742]
TickLength = [0.01 0.025]
TickDir = in
TickDirMode = auto
TightInset = [0.0285714 0.0309524 0.0142857 0.0142857]
Title = [397.002]
Units = normalized
View = [0 90]
XColor = [0 0 0]
XDir = normal
XGrid = off
XLabel = [394.002]
XAxisLocation = bottom
XLim = [0 10]
XLimMode = auto
XMinorGrid = off
XMinorTick = off
XScale = linear
XTick = [0 5 10]
XTickLabel =
    0
    5
    10
XTickLabelMode = auto
XTickMode = auto
YColor = [0 0 0]
YDir = normal
YGrid = off
YLabel = [395.002]

```

```
YAxisLocation = left
YLim = [0 15]
YLimMode = auto
YMinorGrid = off
YMinorTick = off
YScale = linear
YTick = [0 5 10 15]
YTickLabel =
    0
    5
   10
   15
YTickLabelMode = auto
YTickMode = auto
ZColor = [0 0 0]
ZDir = normal
ZGrid = off
ZLabel = [396.002]
ZLim = [-1 1]
ZLimMode = auto
ZMinorGrid = off
ZMinorTick = off
ZScale = linear
ZTick = [-1 0 1]
ZTickLabel =
ZTickLabelMode = auto
ZTickMode = auto

BeingDeleted = off
ButtonDownFcn =
Children = [ (2 by 1) double array]
Clipping = on
CreateFcn =
DeleteFcn =
BusyAction = queue
HandleVisibility = on
HitTest = on
Interruptible = on
Parent = [16]
Selected = off
SelectionHighlight = on
Tag =
Type = axes
UIContextMenu = []
UserData = []
Visible = on
```


Seeing this long list shows what a wealth of options are associated with `plot`. Looking through the list, you see some terms you have already encountered, such as `xlim` and `ylim`, but many new ones as well.

To illustrate how you can make use of the properties in this list, the next program shows how you can control the tick marks in a graph. You can do this using the `set` function. `set` is a very important function because it can be used flexibly in connection with any object property of interest, such as the axes of the current figure. Code 9.12.2 exploits this capability by indicating that the x-axis tick marks run from 2 to 24 in increments of 2.

Code 9.12.2:

```
figure(17)
x = linspace(0,4*(2*pi),100);
y = sin(x);
plot(x,y);
plot(x,y,'g-');
hold on
x_offset = 0;
y_offset = .2;
axis([min(x)-x_offset, max(x)+x_offset,...
      min(y)-y_offset, max(y)+y_offset]);
plot(x,y,'o','color','r','markersize',6,...
      'markeredgecolor','k','markerfacecolor','r');
xlabel('Time');
ylabel('Happiness');
title('Life has its ups and downs. ');
set(gca,'xtick',[2:2:24]);
shg
```

Output 9.12.2:




Here is another example, in which tick marks and associated numbers are suppressed entirely. This can be useful when you want to show a qualitative relation. (The first author

often shows such graphs in his undergraduate teaching when he wants to expose students to a general data pattern.) Only the last line, concerning `xtick`, has been changed from Code 9.12.3, and a new line, concerning `ytick`, has been added.

Code 9.12.3:

```
figure(18)
plot(x,y,'g-');
hold on
x_offset = 0;
y_offset = .2;
axis([min(x)-x_offset, max(x)+x_offset,...
      min(y)-y_offset, max(y)+y_offset]);
plot(x,y,'o','color','r','markersize',6,...
      'markeredgecolor','k','markerfacecolor','r');
xlabel('Time');
ylabel('Happiness');
title('Life has its ups and downs.');
```

Output 9.12.3:



9.13 Plotting Data Points With Error Bars

It is often desirable to show how variable data are by including error bars. These bars often extend above and below a depicted mean by an amount equal to the standard deviation, standard error, or some other measure of variability for the associated sample.

Code 9.13.1 shows how you can display error bars using MATLAB's `errorbar` command. This command takes three arguments: the horizontal position of each point (`x`), the vertical position of each point (`y`), and the length of the bar (`sd`). As shown in Code 9.13.1, the color of the bars and lines can be indicated as well. Here we request black (`'k'`) bars. When the `errorbar` function is used, it tends to connect successive data points with

lines. To hide these lines, you can have MATLAB connect successive data points with white lines ('w-'), as in the code below.


In the first call to `errorbar`, below, only one value of `sd` is specified, so the error bars are the same size above and below each point. A line is specified.

In the second call to `errorbar`, two values of `sd` are specified. The array `sddown`, which has the same length as `sdup`, is all zeros, so only the upward error bars are visible. The lower ones, being zero, are drawn but are invisible.

Code 9.13.1:

```
figure(19)
x= [1:10];
y1 = [4 11 25 65 141 191 313 301 487 673];
sd = [20 30 40 50 58 69 82 78 42 62];
box on
hold on
errorbar(x,y1,sd,'ko-','markersize',6)
hold on
y2 = 700-y1;
sdup = sd;
sddown = zeros(length(sdup),1);
errorbar(x,y2,sddown,sdup,'k.','markersize',18)
shg
```

Output 9.13.1:



9.14 Generating Polar and Compass Plots


So far we have only plotted data in Cartesian coordinates (i.e., rectilinear frames of reference). For data that can be characterized in terms of an *angle* and a *magnitude*, it is possible to plot the data in *polar* coordinates. In these so-called polar plots, each point is positioned some distance (or magnitude) away from the origin along a line with a specified angle relative to the line extending from the origin to the right. The `polar` command allows plotting

points anywhere in this space. A related command, `compass`, draws vectors starting from the origin. In this example, we make a `compass` plot of the three vectors discussed in Section 4.8.

Code 9.14.1:

```
originalradiuspoints = [1 0];
radiusrotated30deg_from_original = [.866 .5];
radiusrotated150deg_from_original = [ -.866 .5];
h1 = compass(originalradiuspoints(1),...
             originalradiuspoints(2));
hold on;
set(h1,'linestyle','-','linewidth',3);
h2 = compass(radiusrotated30deg_from_original(1),...
             radiusrotated30deg_from_original(2));
set(h2,'linestyle','--','linewidth',3);
h3 = compass(radiusrotated150deg_from_original(1),...
             radiusrotated150deg_from_original(2));
set(h3,'linestyle',':','linewidth',3);
```

Output 9.14.1:



9.15 Generating Histograms

Another kind of graph supported by MATLAB is the histogram. A histogram shows the number of elements in various data bins.

Code 9.15.1 shows how to generate a histogram using the `hist` command. The random number generator is first initialized to the default value. Then a 1×2000 matrix of normally distributed random numbers is centered around 5, using `randn`. The `hist` function has two input arguments: the sample to be plotted, and the midpoints of the bins. Here, we specify seven midpoints. If the second argument is omitted, `hist` will make 10 bins by default.


`hist` returns two outputs. One is `N`, a matrix whose elements are the number of values in each of the seven bins that `hist` creates by default. The other is `X`, a matrix whose

elements are the means of the values in each of the seven bins. When `hist` is called again with no explicit outputs, it yields a graph. The colors of the bars in the graph can be set to gray via the command `colormap([.5 .5 .5])`. These numbers signify that in this particular case the values of red, green, and blue are all `.5`. The bars can be brightened by, say, 75% using the command `brighten(.75)`.

Code 9.15.1:

```
figure(21)
rng('default')
sample = randn(1,2000) + 5;
[N,X] = hist(sample,[2:8])
hist(sample,[2:8])
colormap([.5 .5 .5])
brighten(.75)
```

Output 9.15.1:



Output 9.15.2:


N =	6	125	473	788	476	108	24
X =	2	3	4	5	6	7	8

9.16 Generating Bar Graphs

Histograms are just one kind of bar graph. Another kind can be obtained via Code 9.16.1. Here we generate horizontal bars using the `barh` function. (Vertical bars are generated with `bar`.) The bars are gray, as in the last example, but other colors are possible, as indicated in the comments concerning `colormap`. In the code below, the value assigned to `brighten` is smaller than before and the bars are, accordingly, darker than in the last output.

Code 9.16.1:


```
figure(22)
a = [3 4 5 6 7 6 5 4 3];
barh(a)
colormap([.5 .5 .5]) % gray bars
% colormap([0 0 0]) % black bars
% colormap([1 1 1]) % white bars
% colormap([1 0 0]) % red bars
% colormap([0 1 0]) % green bars
% colormap([0 0 1]) % blue bars
brighten(.15)
ylim([0 ,10])
xlim([0 8])
```

Output 9.16.1:

As the comments in the above code indicate, it is possible to plot bars in different colors using the three values of `colormap`. When all three values are the same, it is possible to generate grayscale values:

Code 9.16.2:

```
close all;
clear all;
figure(22)
data = [
1 2 3
4 5 6
7 8 9];
bar(data);
colors = [
1 1 1
.5 .5 .5
0 0 0];
colormap(colors)
```

Output 9.16.2:**9.17 Saving, Exporting, and Printing Figures**

How can you export and print figures from MATLAB? The simplest method is to manually copy one figure at a time and paste it into the document you're producing (e.g., a Word document for a grant proposal you're writing). To do this, keep the figure window open, click on the Edit icon of the toolbar, and then click on Copy Figure.

You can also save figures from the Command window or your program code using the `saveas` command. By default (if you specify no other file extension) figures are saved in native MATLAB format as `.fig` files. These can be opened in MATLAB using `load`, which restores the file just as it looked when the program ran, with all the image tools at the top of the figure window again available. Other ways of saving figures are needed for working with other programs. If you are working with Figure 5, for example, `saveas(5, 'mysinwaves.jpg')` will save it to a file called `mysinwaves` using the `.jpg` format; `saveas(5, 'mysinwaves.tif')` will save the file as a `.tiff` file. Each of these files can later be loaded into a MATLAB program using the `imread` command (see Section 10.3) or opened with other graphic or word-processing programs.

A third, more flexible, method is to use the `print` command, as in the following examples. The `print` command does not actually print your figure on paper, but rather generates a printable file that can be printed using another graphics program. In the first example, the current figure, figure (23) is saved at 600 pixel resolution (`-r600`) to a `.jpeg` file (`-djpeg`) named `Figure_9_17_1.jpg`. In the second example, figure (24) is saved to a `.tif` file (`-dtiff`) named `Figure_9_17_2.tif`. In the third example, figure (25) is saved to an `.eps` file (`-deps`) named `Figure_9_17_3.eps`. The `loose` option ensures that there is a border around the graph; otherwise the graph expands like a balloon to fill the available space. By default, the files are created in the current working directory.

You can check that the files have been created by using the `ls` command, though your validation of the figures is complete only when have opened, inspected, and approved them. The fine details of the files differ, as can be seen by comparing the detail in the numeral '1' at the top left of each figure. Note that the `.jpg` file shows some smoothing at high resolution, but also some artifacts due to compression of the image. The `.tif` file captures exactly what is on the screen (pixel-by-pixel) and so it is quite jagged (as the screen would be if you looked

closely enough). The `.eps` file has the clearest resolution, regardless of magnification, and is, for this kind of graphic, usually the preferred file type for publication. The figures used for this text were, for the most part, generated as `.eps` files. However, unlike the `.fig`, `.jpg`, or `.tif` files, the `.eps` type file cannot be read into MATLAB again.


Code 9.17.1:

```
figure(23)
plot([1:10],[1:10].^-2,'k-o')
print -r600 -djpeg Output_9_17_1


figure(24)
plot([1:10],[1:10].^-2,'k-s')
print -dtiff Output_9_17_2


figure(25)
plot([1:10],[1:10].^-2,'k-^')
print -deps -loose Output_9_17_3
```

Output 9.17.1 (detail, *.jpg* format file):



Output 9.17.2 (detail, *.tif* format file):



Output 9.17.3 (detail, .eps format file):

There are several alternative printing formats (`-depsc`, `-depsc2`, `-dpdf`, etc.) described in the MATLAB documentation that might be best for certain graphics (e.g. color). For image files, the `-r600 -djpeg` version may be best. One other useful point to keep in mind is that if your program changes the size of the figure to be printed from the default figure size, the output file may be distorted. The command, `set(gcf, 'paperpositionmode', 'auto')` just before the print command will often rectify this problem.

9.18 Generating Other Kinds of Graphs and Getting and Setting Figure Properties

It is worth considering a few summary statements at this point. First, most of the techniques described earlier for regular line graphs also apply to bars and histograms. For example, you can use `xlabel`, `ylabel`, and `title` with bars and histograms. The best way to see what works is to experiment!

Second, there are other plot options that you can explore for yourself. If you want to have different coordinates on the left and right vertical axes, you can use `plotyy`. MATLAB also lets you make special plots based on the `stairs` command, the `stem` command, the `pie` command, the `feather` command, and the `quiver` command. You should know enough plotting from this chapter to explore these other options on your own.

Third, `get(gcf)` gets you properties of the current figure using `gcf`, which stands for “get current figure,” and is a shortcut for the handle of the active window. You can also use `gcf` to set other properties of interest to your taste, just as we did in Section 9.10 for the abscissa label. Typically in MATLAB, if you can get a property of any object (figure, axis, text, or line), you can set it to a new value, though a few properties are “read-only” and so are unmodifiable. Here is an example of how you can control the size of a figure.

Code 9.18.1:

```
| set(gcf, 'Position', [100 200 500 500])
```

The values in the array are, respectively, the coordinates of the left and bottom edges of the figure (relative to the bottom left of the screen) and the figure’s width and height. You can

find values you like by manually repositioning and resizing a figure until you like it, then use `get(gcf, 'Position')` to determine what those pleasing values are. Finally, you can enter those values into a line of code like Code 9.18.1. If you work on different computers, you can use `get(0, 'Screensize')` to see the dimensions of your screen before you resize the window.

A final remark is that this chapter has only scratched the surface of things that can be done with plots in MATLAB. Because the aim of this book is to equip you with the intellectual tools needed to get you started with this programming language, you should know enough from this chapter to create your own two-dimensional graphs and draw on the wealth of information in MATLAB's Help documents and related sources to see for yourself how "the plot thickens."

9.19 Practicing Plots

Try your hand at the following exercises, using only the methods introduced so far in this book or in information given in the problems themselves.


Problem 9.19.1:

The following code will yield one bell-shaped curve. Modify the code to get two bell-shaped curves, with one shifted .5 units to the right of the other, as shown in Output 9.19.1.

Code 9.19.1:

```
figure(1)
x = linspace(0,1,200);
a = 6;
b = 6;
y = (x.^a) .* ((1-x).^b);
plot(x,y,'k')
```

Output 9.19.1:




Problem 9.19.2:


Problem 5.8.5 referred to the equation

$$p_correct = base_rate + learning_rate * \log(trial),$$

where `trial` could take on the values 1, 2, 3, . . . , 200, `learning_rate` could be any real number between 0 and 1, `base_rate` was .25, and `p_correct` could not exceed 1. Generate a figure resembling the one below by setting `learning_rate` to .02. Plot `p_correct` as a function of `trial`, label the x axis `Trials`, label the y axis `Proportion Correct`, and have the title say `Learning`. Have the points appear as black `o`'s connected with line segments. `grid` should be on, `box` should be on, and the plot should appear in figure (2) .

Output 9.19.2:**Problem 9.19.3:**


Adapt the program you wrote for the last problem to generate a figure resembling the one below by setting `learning_rate` to .02, .04, and .06. Have the plot appear in figure (3) .

Output 9.19.3:

Problem 9.19.4:

Adapt the program you wrote for the last problem to generate a figure resembling the one below by again setting `learning_rate` to `.02`, `.04`, and `.06` and making the subplots on the right show the cumulative number correct for each of the three learning rates. Have the plot appear in figure (4).


Output 9.19.4:



Problem 9.19.5:

Adapt the program you wrote for the last problem to generate a figure that resembles the one below. There are two new features of the figure to be generated. One is that the learning rates are specified as text in each of the left subplots. The other is that the subplots on the right include a star at the trial for which the cumulative number correct exceeds 50. Have the plot appear in figure (5).

Output 9.19.5:



Problem 9.19.6:

Consider these made-up data:

Lefthanders:

Condition	RT (ms)
Valid	
Left	240
Right	230
Invalid	
Left	270
Right	260

Righthanders:

Condition	RT (ms)
Valid	
Left	210
Right	220
Invalid	
Left	280
Right	290

Plot the data in two adjacent 2×2 subplots with appropriately labeled axes. Try bar graph and line graph styles. Make nice big points. By inspection, does there seem to be a statistical interaction in this hypothetical experiment? Which kind of graph shows this most convincingly?

10. Lines, Shapes, and Images

This chapter covers the following topics:

- 10.1 Generating lines
- 10.2 Forming and filling shapes
- 10.3 Loading images
- 10.4 Generating your own images
- 10.5 Clicking in figure windows to add graphics, add text, or record responses
- 10.6 “Stairing”
- 10.7 Generating three-dimensional bar graphs
- 10.8 Plotting in three dimensions
- 10.9 Plotting above a meshgrid
- 10.10 Plotting “meshy” data
- 10.11 “Surfing” the “web”
- 10.12 Changing points of view
- 10.13 Generating contours
- 10.14 Checking your understanding of meshgrid-based graphing
- 10.15 Generating rectangular solids
- 10.16 Generating spheres and cylinders
- 10.17 Generating ellipsoids
- 10.18 Practicing plots

The commands that are introduced and the sections in which they are premiered are:

<code>set(gca)</code>	(10.1)
<code>set(gcf)</code>	(10.1)
<code>axis square</code>	(10.2)
<code>fill</code>	(10.2)
<code>image</code>	(10.3)
<code>imread</code>	(10.3)
<code>axis equal</code>	(10.4)
<code>colormap</code>	(10.4)
<code>fontsize</code>	(10.5)
<code>ginput</code>	(10.5)
<code>rotation</code>	(10.5)
<code>stairs</code>	(10.6)
<code>bar3</code>	(10.7)

<code>plot3</code>	(10.8)
<code>zlabel</code>	(10.8)
<code>meshgrid</code>	(10.9)
<code>mesh</code>	(10.10)
<code>surf</code>	(10.11)
<code>view</code>	(10.12)
<code>contour</code>	(10.13)
<code>surfc</code>	(10.14)
<code>surfl</code>	(10.14)
<code>zlim</code>	(10.14)
<code>patch</code>	(10.15)
<code>cylinder</code>	(10.16)
<code>sphere</code>	(10.16)
<code>camtarget</code>	(10.17)
<code>camzoom</code>	(10.17)
<code>light</code>	(10.17)
<code>rotate</code>	(10.17)
<code>shading</code>	(10.17)

10.1 Generating Lines


In the last chapter, you learned about data plots, and you were exposed, only in passing, to lines, shapes, and images. Those elements need not be used only in data plots, however. They can also be used for other purposes. By gaining greater mastery of these graphic elements, you can enhance the figures you generate, whether in data plots or in other contexts.

Begin with lines. You generated lines with the `plot` command (e.g., Output 9.8.1). To supplement that material, consider Code 10.1.1. Here we `clear` all variables, `close` all figure windows, and provide instructions for drawing a line in figure (1). We use the `plot` command, recalling that this command takes two arguments: an `x` (abscissa) array and a `y` (ordinate) array. In this instance, we limit the `x` array to just two values, the starting and ending values of `x`. We also limit the `y` array to two values, the starting and ending values of `y`. For aesthetic reasons, we enclose the graph in a box, using the `box on` command. We assign `plot(x, y)` to a variable called `our_first_line` so we can later use this handle to manipulate the line.

Code 10.1.1:

```
clear all
close all

figure(1)
x = [0 1];
y = [0 1];
box on
our_first_line = plot(x, y);
```

Output 10.1.1:

We can examine the properties of `our_first_line` by calling the `get` function.

Code 10.1.2:

```
get(our_first_line)
```

Output 10.1.2:

```
Color = [0 0 1]
EraseMode = normal
LineStyle = -
LineWidth = [0.5]
Marker = none
MarkerSize = [6]
MarkerEdgeColor = auto
MarkerFaceColor = none
XData = [0 1]
YData = [0 1]
ZData = []

BeingDeleted = off
ButtonDownFcn =
Children = []
```

```

Clipping = on
CreateFcn =
DeleteFcn =
BusyAction = queue
HandleVisibility = on
HitTest = on
Interruptible = on
Parent = [151.008]
Selected = off
SelectionHighlight = on
Tag =
Type = line
UIContextMenu = []
UserData = []
Visible = on

```

Having discovered that a property of `our_first_line` is `color`, we can specify the `color` for a new plot. The new plot will be displayed in figure (2) and will be assigned to the variable `our_second_line`. We have to issue the `box on` command again if we want the box to be on. The reason is that `box` is set to `off` each time a new figure window is opened.

By saying `'color', [1 0 0]`, we indicate that we want the value of red to be 1 and we want the values of green and blue to both be 0. Remember that the three numbers in `'color'` matrix are the proportions of red, green, and blue. In the last example, the line was blue, as indicated by the first line of Output 10.1.2, `Color = [0 0 1]`. Although the line in this book is black, the actual, intended color can be seen on this book's website (www.routledge.com/9780415535946), or you can run Code 10.1.3 and see it on your monitor.


Code 10.1.3:

```

figure(2)
delta_y = .5;
our_second_line = plot([min(x) max(x)],...
                      [min(y)+ delta_y max(y) + delta_y], 'color', [1 0 0]);
box on

```

Output 10.1.3:




Having discovered that another property of `our_first_line` is `linestyle`, we can specify a new `linestyle` and, for that matter, a new `color` and `linewidth`. In Code 10.1.4, we specify these values for `our_third_line`, to be drawn in figure (3). We use the `set` command rather than the `get` command this time around because we are using the handle to set one or more of its properties. We also experiment with new values of `color` so the red, green, and blue elements of the `color` matrix are not just assigned 1's and 0's. The color values used below make for a bright brown, as can be seen on a monitor if you run the program.

Code 10.1.4:

```
figure(3)
delta_y = 1;
our_third_line = plot([min(x) max(x)],...
    [min(y)+ 2*delta_y max(y) + 2*delta_y]);
set(our_third_line,'color',[.9 .5 .1], ...
    'linestyle','--', ...
    'linewidth',8);
box on
```

Output 10.1.4:



If you don't want all the properties of a `plot` but instead want specific properties, you can type `set(gcf)` to find out about figure properties, or you can `set(gca)` to find out about axis properties. You can find out about specific figure properties or axis properties by adding optional strings to inquire about them, as in these two examples:

Code 10.1.5:

```
set(gca,'XGrid')
set(gcf,'PaperOrientation')
```

Output 10.1.5:

```
[ on | {off} ]
[ {portrait} | landscape | rotated ]
```

Output 10.1.5 gives the possible values of the properties in question. The values in brackets are the default values—what MATLAB provides when no specific, alternative instructions are supplied.

10.2 Forming and Filling Shapes

Shapes are enclosed n -sided polygons, where $n \geq 3$ is the number of straight line segments enclosing the polygon. Thus, a triangle is an $n = 3$ shape, a rectangle is an $n = 4$ shape, and so on. When the lengths of the straight line segments are equal, the $n = 3$ shape is an equilateral triangle and the $n = 4$ shape is a square.

MATLAB provides a function called `fill`, which lets you form and fill shapes. We use the `fill` function in Code 10.2.1 in a function called `my_polygon_1`. This function takes three arguments. The first is the number of sides, n , of the polygon to be filled. The second is the distance, r , of each vertex of the polygon from the polygon's center. (r is effectively the radius of a circle when n is so large that the generated shape is visually indistinguishable from a circle.) The third argument is the 3×1 RGB color matrix defining the color. The first number defines the proportion of red, the second number defines the proportion of green, and the third number defines the proportion of blue. `my_polygon_1` uses an x matrix and a y matrix as well as some trigonometry (see Sections 9.1 and 9.4). We add 1 to n because $n + 1$ vertices must be specified to generate n sides; the polygon's enclosing line must return to its origin. The call to `my_polygon_1` is shown in Code 10.2.2, where we indicate that in figure (4), we wish to fill a four-sided polygon whose “radius” has length 1, and whose color is given by the matrix `[.5 .5 .5]`. We also use `axis square` to keep the current axes the same size. For other axis options, we know we can turn to `help axis`.

Code 10.2.1:

```
function my_polygon_1(n,r,c)


x = linspace(0,2*pi,n+1)
x = r*cos(x)

y = linspace(0,2*pi,n+1)
y = r*sin(y);

fill(x,y,c)
```

Code 10.2.2:

```
figure(4)
my_polygon_1(4,1,[.5 .5 .5])
axis square
```


Output 10.2.2:

The four-sided polygon in Output 10.2.2 is a diamond. If you want a square, you need to rotate the shape. In the function `my_polygon_2`, introduced in Code 10.2.3, we add a fourth argument that provides for such rotation. The fourth argument is called `turn`. It shifts the series of angles used to define `x` and `y`. The new computation that uses `turn` is designed so you can set the fourth term to 0 for the default orientation. Calls to `my_polygon_2` are shown in Code 10.2.4, where the angles increase from a negative value up to 0, so a square is the last polygon drawn, making it the one that sits “on top” of the others.

Code 10.2.3:


```
function my_polygon_2(n,r,c,turn)
x = linspace(0,2*pi,n+1)
x = x + (turn + 1/(2*n))*(2*pi);
x = r*cos(x);

y = linspace(0,2*pi,n+1)
y = y + (turn + 1/(2*n))*(2*pi);
y = r*sin(y);

fill(x,y,c)
```

Code 10.2.4:


```
figure(5)
hold on
for turn = linspace(-.2,0,5)
    my_polygon_2(4,1,[.5 .5 .5],turn)
    axis off
end
```

Output 10.2.4:

The `fill` command applies to irregular polygons and even to closed forms whose line segments cross, as illustrated in Code 10.2.5. Here, a “crazy” series of `x` and `y` values are used. In addition, a property of the object being drawn—the width of the edge line—is specified through the `set` command.

Code 10.2.5:

```
figure(6)
crazy_x = rand(1,5);
crazy_y = rand(1,5);
f = fill(crazy_x,crazy_y,'g')
set(f,'LineWidth', 5.0);
```

Output 10.2.5:**10.3 Loading Images**


All the figures considered so far have been generated with MATLAB code, but MATLAB also permits loading of images from other sources. Below, we load an image that was saved

earlier in .jpg format. Two new commands are used. One is `imread`, which takes as its argument the name of the file to be loaded, enclosed in quote marks (needed because it is a string) along with its file type (.jpg). Note that a semi-colon appears at the end of the line containing `imread`. This semi-colon is *extremely* important. Without it, the Command window would show a flood of numbers, reflecting the vast amount of information contained in an image, even in a relatively simple one like the photograph displayed in Output 10.3.1. The image shown in Output 10.3.1 is presented via the `image` command, which takes as its argument the variable created with `imread`. In Code 10.3.1 we turn off the axis for aesthetic reasons. The photograph was taken by one of the authors.

Code 10.3.1:

```
figure(8)
a = imread('view_from_window.jpg');
image(a)
axis off
```


Output 10.3.1:



Here is another example showing that photographs can be displayed via the procedures shown above. The picture shows the first author testing a participant in a study of perceptual-motor control. The participant, who agreed to let his image be shown here, takes hold of a plunger to transport it to the platform to the right. The datum of interest is where the plunger is grasped as a function of the height of the target platform. The main finding is that grasp heights are inversely related to target heights (Cohen & Rosenbaum, 2004).

Code 10.3.2:

```
figure(9)
b = imread('lab_photo.jpg');
image(b)
axis off
```


Output 10.3.2:**10.4 Generating Your Own Images**

Because an image is represented as a two-dimensional matrix of pixel values, you can generate your own image with MATLAB. There are two ways to represent a color image of size x pixels by y pixels. The first is as an $x \times y \times 3$ matrix, where the color of the pixel at each of the two-dimensional (x, y) points in the figure is represented by the combination of third-dimension values.

Think of an image as a matrix with three sheets, one on top of the other. The first layer is the intensity of *red* at each of the sheet's horizontal and vertical positions, the second layer is the intensity of *green* at each of the sheet's horizontal and vertical positions, and the third layer is the intensity of *blue* at each of the sheet's horizontal and vertical positions. Judicious selection of values in each of the three sheets for each horizontal and vertical position lets you define the color for that position. Here is an example. (Run the program or see the website to appreciate the colors).

Code 10.4.1:

```
clf;
clear;
Rainbow(1,1:6,1:3) = [
    1 0 0    % Red
    1 .5 0   % Orange
    1 1 0    % Yellow
    0 1 0    % Green
    0 0 1    % Blue
    1 0 1    % Violet
];
image(Rainbow)
axis equal
```

Output 10.4.1:

The other way of representing an image is to use a color map. In this case, the image is represented as an $n \times n$ matrix rather than a full $n \times n \times 3$ matrix. Each cell in the matrix contains an integer that stands for one of the c different colors used in the image. The specific color of each pixel is determined by using that integer, c , as the index to a color map, a $c \times 3$ matrix of colors that maps those integers to different RGB color values.

Here is an example in which we use a color map for which, arbitrarily, 1 = white and 2 = black, so the color map is a 2×3 matrix (two colors $\times$ the three RGB values, one for each color). The relation between the integer value in the matrix of the image and the color represented in the color map is arbitrary. The cell's value simply serves as an index into the color map.

To make the example interesting, we next use it to generate a random-dot stereogram with separate views for each of the two eyes. We impose some retinal disparity between the two images, so when the two images are superimposed, the two images can trigger stereoscopic depth perception (Julesz, 1971; Rock, 1985). If you fuse the images by convergent eye movements, you will see the central square pop out in depth. If you fuse the images via divergent eye movements (looking “beyond” the page) the square will recede. We use `subplot` to make the separate images as two axes in the same figure.

Here are the mechanics of the program. We start by positioning the window on the monitor (see Section 9.18). The `Make Identical Images` section assigns a value of 1 or 2 to each cell of a 40×40 matrix, using the `randi` command. The `Superimpose` an inner square and `shift` section puts a second random 20×20 matrix in the middle of the larger squares, shifted one pixel to the left or right in each eye. The color map is defined as a matrix with the first row `[1 1 1]`, representing white, and the second row `[0 0 0]`, representing black, so if a pixel of the matrix has a 1, it will be shown in white, whereas if it has a 2, it will be shown in black. If you now fuse the two images by converging (“crossing”) your eyes, you may see a smaller square hovering in front of the larger one, because the two eyes are receiving slightly different patterns in the images, corresponding to the cues for stereoscopic depth.

Code 10.4.2:

```
% Set up
figure(10)
windowposition = [10,550,1000 500];
set(10,'Position',windowposition);
% Make Identical Images
RightEyePicture = randi(2,40,40);
LeftEyePicture = RightEyePicture;
% Superimpose and shift an inner square
InnerSquare = randi(2,20,20);
RightEyePicture(11:30,11:30) = InnerSquare;
LeftEyePicture(11:30,13:32) = InnerSquare;
% Define color map, row 1 = white, row 2 = black
mycolors = [
    1 1 1
    0 0 0
];
colormap(mycolors);
% Display
subplot(1,2,1);
image(RightEyePicture);
title('Right','fontsize',16)
axis off;
axis equal
subplot(1,2,2);
image(LeftEyePicture);
title('Left','fontsize',16)
axis off;
axis equal
shg
```

Output 10.4.2:

To explore this further, you can change the values in `mycolors` to vary the contrast between the pixels. For a similar example of the use of an $n \times n$ image array that maps to a 10×3 color map defining 10 colors, see the code used in the following entertaining demonstration featuring the cartoon character, Homer Simpson: (www.mathworks.com/matlabcentral/fileexchange/12079-forbidden-donut/content/fdonut.m).

10.5 Clicking in Figure Windows to Add Graphics, Add Text, or Record Responses

The photograph in Output 10.3.2 includes stuff that is neither particularly relevant to the study nor especially pretty. It would be nice to hide the section of shelf with the tape measure, folders, and glasses. We use this rather mundane challenge as a way of introducing a useful capability of MATLAB, namely, recording where someone clicks in a figure window. On the basis of this clicked information, it is possible to add graphics, add text, or record responses. In the present example, we want to add graphics that covers the unsightly junk.


One command that makes such things possible is `ginput`. This command is used in Code 10.5.1 in connection with `figure(9)`, which was shown in Output 10.3.2. With `figure(9)` active and with `hold on`, `ginput(2)` tells MATLAB to expect two clicks in the current figure window. More generally, `ginput(n)` tells MATLAB to expect $n \geq 1$ clicks in the current figure window. `ginput` by itself (with no argument supplied) tells MATLAB to expect an indeterminate number of clicks, until the return (Enter) key is pressed. See `help ginput` for more information about this very useful command.

When `ginput` is called, crosshairs appear in the figure window where the mouse is currently positioned. Moving the mouse causes the crosshairs to move. When the crosshairs are in a desired position, you can click the mouse and the crosshairs' (x , y) coordinates will be recorded.

In Code 10.5.1, just two clicks are collected because we want to cover the extraneous part of the image with a rectangle, only two of whose corners—the bottom left and top right, or the top left and bottom right—need to be clicked for the rectangle to be defined. The two values of x and the two values of y are collected in `[x y] = ginput(2)`. The two values of x and y can define the four corners of a rectangle we will draw using the `fill` command. We will make the added rectangle white to have blend it in with the white of the page.

Code 10.5.1:


```
b = imread('lab_photo.jpg');  
image(b);  
hold on;  
[x y] = ginput(2);  
xs = [x(1) x(2) x(2) x(1)];  
ys = [y(1) y(1) y(2) y(2)];  
fill(xs,ys,'w');
```

Output 10.5.1:

The `ginput` command also makes it possible to add text to an active figure window at a clicked location. In Code 10.5.2, we collect just one click at a location where we want the first character of the text to be drawn. We tell MATLAB to draw text that has two properties, not previously introduced in this book. One is `'rotation'`, which is here set to 90 degrees. When `'rotation'` is not specified, its default value is 0 degrees. The second property is `'fontsize'` which here is set to 24 point. When `'fontsize'` is not specified, its default value is 12.

Code 10.5.2:

```
clear x y
[x y] = ginput(1);
text(x,y,'Take the plunge!','rotation',90,'fontsize',24);
```

Output 10.5.2:**10.6 “Stairing”**


The `stairs` command can be used to plot data where vertical line segments are joined at their ends by horizontal line segments, all going to the right or left, and where horizontal

line segments are joined at their ends by vertical line segments, all going up or down. Such graphs are often used in psychophysics to describe the staircase method of tracking a threshold. Similarly, in operant conditioning research, continuous performance of a subject (typically, a rat, pigeon, or human) is conventionally presented in the form of a *cumulative record* (Skinner, 1972), which resembles a flight of stairs. The cumulative record shows the total number of responses accumulated over time. If the responses are lever presses made by a rat, each press steps up a line by one unit, whereupon the line extends to the right as time passes until the next press is made, at which point the line steps up another unit. In the standard cumulative record, responses that lead to reinforcement are indicated by a marker.

To demonstrate the cumulative record with hypothetical data, we generate some pseudo-data for 30 time intervals by assigning a zero or 1 response to each interval (with no responses after interval 20), and then use the `cumsum` function to accumulate the responses; `cumsum` yields a cumulative sum. (Our reference to pseudo-data calls to mind ethical issues, which we will take up at the end of the book.) A random half of the responses are then identified as having earned reinforcement. Finally, the cumulative response record is plotted, and the reinforcement marks are added for those responses that were reinforced.

Code 10.6.1:

```
responses(1) = 0;
responses(2:21) = randi(2,1,20)-1;
    responses(22:30) = 0;
cumrec = cumsum(responses);
reinforcedtrials = [];
for i = 1:20
    if responses(i) > 0 && randi(2) > 1
        reinforcedtrials = [reinforcedtrials i];
    end
end
stairs(cumsum(responses));
hold on
for i = 1:length(reinforcedtrials)
    j = reinforcedtrials(i);
    plot([j j+.5], [cumrec(j) cumrec(j)-.5]);
end
xlabel('Time','FontSize',16);
ylabel('Cumulative Responses','FontSize',16);
```


Output 10.6.1:**10.7 Generating Three-Dimensional Bar Graphs**

It is often useful to visualize data in three dimensions, especially when the data define or describe three-dimensional objects. MATLAB has methods for such visualization.

We begin our discussion of 3-D visualization with the extension of simple two-dimensional bar graphs to three dimensions. In the code that follows, we consider hypothetical frequency histograms corresponding to normally distributed samples of different sizes. We plot the histograms as a set of ordinary bar graphs and then plot the histograms in three-space using the `bar3` command. The main point of the example is that single three-dimensional graphs can give a different (often better) perspective than can multiple two-dimensional graphs. As before, see the book's website or run the program to see the figure in color.

Code 10.7.1:

```
m = [];
for j = 2:4
    clear n x
    [n x] = hist(randn(1,10^j),10);
    subplot(3,2,((j-1)*2)-1)
    bar(n)
    m = [m;n];
end
subplot(3,2,[2 4 6])
bar3(m')
```

Output 10.7.1:


10.8 Plotting in Three Dimensions

The next example shows how data can be plotted in three dimensions with the `plot3` command. The example is based on one provided in MATLAB's Help regarding `plot3`, although `xlabel`, `ylabel`, `zlabel`, and `title` have been added in the code that follows. `zlabel` is used here for the first time. In this example, `sin(t)` and `cos(t)` are used to describe a circular function (in the x-y axes). When this function is plotted in three dimensions with `t` in the z-axis, the result is a spiral.

Code 10.8.1:

```
figure(3)
t = 0:pi/50:10*pi;
plot3(sin(t),cos(t),t)
axis square;
grid on
box on
xlabel('sin(time)', 'rotation', 0);
ylabel('cos(time)', 'rotation', 0);
zlabel('time');

title('Slinky');
```


Output 10.8.1:**10.9 Plotting Above a Meshgrid**

The graph in Output 10.8.1 can be viewed as a trajectory—for example, of a hawk spiraling upward in an updraft. A trajectory has the property that there is only one way to get from one point to another within the plotted object. Sometimes, however, you want to look at an entire surface, the defining feature of which is that there are many possible paths between points. Here the `meshgrid` function is useful. `meshgrid` creates a grid of values forming a mesh on the “floor,” with values corresponding to each intersection point plotted “above” it.

The following code, which is slightly adapted from MATLAB’s help about `meshgrid`, shows how `meshgrid` can be used. In this example, a matrix, $[X, Y]$, is created with the `meshgrid` function applied to a linearly spaced array of 41 elements spanning -2 to 2 . Z values are then plotted “above” the points created with `meshgrid`, in this case according to the equation in the third line of Code 10.9.1. The `plot3` function is then used to plot Z as a function of $[X, Y]$.

Code 10.9.1:

```
figure(4)
[X,Y] = meshgrid(linspace(-2,2,41));
Z = X.*exp(-X.^2 - Y.^2);
plot3(X,Y,Z)
grid on
```


Output 10.9.1:

10.10 Plotting “Meshy” Data

The graph in Output 10.9.1 is a series of disconnected lines. You can connect the lines using the `mesh` command. Additionally, but optionally, you can indicate that you would like your “meshy” data to occupy a box, as in the code below.

Code 10.10.1:

```
| figure(5)
| mesh(X,Y,Z)
| box on
```

Output 10.10.1:


You can regenerate this graph and add points to it with `plot3`.

Code 10.10.2:

```
| figure(6)
| mesh(X,Y,Z)
```

```
hold on
plot3(X,Y,Z, 'k. ')
box on
```

Output 10.10.2:




10.11 “Surfing” the “Web”

The surfaces in Outputs 10.10.1 and 10.10.2 consist of unfilled polygons. It would be desirable to fill the polygons to create a more solid-looking, multi-colored surface. Because a mesh with unfilled polygons looks a bit like a spider’s web, and because the MATLAB command that fills unfilled polygons in a mesh is called `surf`, we have titled this section, partly for amusement, “surfing the web.”

Code 10.11.1 is used to display X , Y , Z using the `surf` command. We give the graph a title (`'Surf's Up!'`) and surround the graph with a box to make it pretty. To learn about properties of the graph’s axes, we write `get(gca)`. As a reminder, if you want to learn about properties of `surf(X,Y,Z)`, you can get `(surf(X,Y,Z))` or `get(s)`, assuming s was previously assigned to `surf(X,Y,Z)` with `s = surf(X,Y,Z)`. Similarly, if you want to learn about properties of the figure window, you can get `(figure(7))` or `get(gcf)`. Note that the graph shown below appears in grayscale. In MATLAB or in the website for this book (www.routledge.com/9780415535946), the peak to the right is in red (signaling positive values) and the peak to the left is in blue (signaling negative values).

Code 10.11.1:

```
figure(7)
surf(X,Y,Z)
title('Surf's Up!')
box on
```

Output 10.11.1:

10.12 Changing Points of View

We had an ulterior motive for getting the axis properties of the graph shown in Output 10.11.1. Apart from the fact that getting such properties helps you identify the properties you can specify, one of the properties was particularly interesting and important, namely `view`.

Code 10.12.1:

```
| help view
```

Output 10.12.1:

```
VIEW    3-D graph viewpoint specification.
VIEW(AZ,EL) and VIEW([AZ,EL]) set the angle of the
view from which an observer sees the current 3-D plot.
AZ is the azimuth or horizontal rotation and EL is the
vertical elevation (both in degrees). Azimuth revolves
about the z-axis, with positive values indicating
counter-clockwise rotation of the viewpoint. Positive
values of elevation correspond to moving above the
object; negative values move below. VIEW([X Y Z]) sets
the view angle in Cartesian coordinates. The magnitude
of vector X,Y,Z is ignored.
```

Because Output 10.11.1 contained the statement,

```
view = [-37.5 30]
```


you can infer that the default values of `view` supplied by MATLAB when `view` is not explicitly specified has an azimuth of -37.5 degrees and an elevation of 30 degrees. Beware that in this context, MATLAB uses degrees, not radians to represent angles.

Suppose you want to look at the data depicted in Output 10.11.1 from directly overhead (i.e., with an elevation of 90 degrees) and, just to keep things simple, at an azimuth of 0 degrees. Relevant code and output follow.

Code 10.12.2:

```
figure(8)
surf(X,Y,Z)
set(gca,'view',[0,90])
```

Output 10.12.2:




The graph in Output 10.12.2 is in grayscale but will appear in color on your monitor if you run the program yourself or go to the website for this book (www.routledge.com/9780415535946). Regions to the left are blue, whereas regions to the right are red. The graph viewed in color looks like graphs often seen in behavioral science. For example, crime rates on different sides of the track are sometimes summarized in graphs like the one shown in Output 10.12.2. Similarly, maps of brain activity are often depicted in terms of more active (red) and less active (blue) regions. If you suppose that the left and right halves of Output 10.12.2 correspond to the left and right hemispheres of the human cerebral cortex, you might surmise that this fMRI (if it were one) came from a task demanding more right- than left-hemisphere activity.

10.13 Generating Contours

Another way to visualize a data pattern like the one shown in Output 10.12.2 is with the `contour` function. This function lets you see “edges” between regions. Contour maps for terrestrial landscapes typically demarcate different height ranges. The contour map below likewise shows demarcations between low levels of z (blue, on the left side) and higher levels of z (red, on the right side).

Code 10.13.1:

```
| figure(9)
| contour(X,Y,Z)
```

Output 10.13.1:

10.14 Checking Your Understanding of Meshgrid-Based Graphing

The foregoing examples concerned alternative ways that three-dimensional data could be graphed when z values were plotted as a function of meshgrid-defined x and y values. To check your understanding of the mapping of z values onto x and y meshgrid-defined values, you can define an $[x, y]$ matrix via `meshgrid(1:8)`, leaving off the semi-colon to see what $[x, y]$ looks like. As shown in Output 10.14.1, $[x, y]$ is actually an x matrix and a y matrix, each of which has size 8×8 in this particular case. y is just the transpose of x . Thinking about this result, you can recall that `meshgrid` generates a distinct x value for every y value, and vice versa.

Code 10.14.1:

```
| [x,y] = meshgrid(1:8)
```

Output 10.14.1:

```
| x =
|
|    1    2    3    4    5    6    7    8
|    1    2    3    4    5    6    7    8
|    1    2    3    4    5    6    7    8
|    1    2    3    4    5    6    7    8
```


1	2	3	4	5	6	7	8
1	2	3	4	5	6	7	8
1	2	3	4	5	6	7	8
1	2	3	4	5	6	7	8
y =							
1	1	1	1	1	1	1	1
2	2	2	2	2	2	2	2
3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	4
5	5	5	5	5	5	5	5
6	6	6	6	6	6	6	6
7	7	7	7	7	7	7	7
8	8	8	8	8	8	8	8

Next, we use the values of x and the values of y to define a z matrix. For the graph we wish to draw, we want the value of z to be small when x is close to the mean of all the x values and to grow quadratically as x departs from the mean of all the x values (see Section 9.9). Similarly, we want the value of z to be small when y is close to the mean of all the y values and to grow quadratically as y departs from the mean of all the y values. Recalling that `mean(x)` will return means for each column of x , and that `mean(y)` will return means for each column of y (see Chapter 3), the line of code in which z is defined uses `mean(mean(x))` and `mean(mean(y))`. We multiply the squared deviations by 10 to make the gradient steeper, and we use `surf` to add “lighting” to the generated surface, and `surf` to show the contours beneath the surface in the final graph, which is here allowed to fill positions 5 and 6 of the 3×2 matrix of subplots. The five graphs use different views based on an initial view generated in exploratory work. We set the limits of the z axis with `zlim`.

Code 10.14.2:

```
figure(10)
z = (10*(x-mean(mean(x))).^2) + (10*(y-mean(mean(y))).^2);


for v = 1:5
    if v < 5
        subplot(3,2,v)
        surf(x,y,z)
    else
        subplot(3,2,5:6)
        surf(x,y,z)
    end
    zlim([0 max(max(z))+2]);
    set(gca,'view',[50.5 v*76.2987]);
end
```

Output 10.14.2:

The code below shows the contour map for the surface.

Code 10.14.3:

```
contour(x,y,z)
```

Output 10.14.3:

In the next program we create a surface with four minima, not just one. A low and a high attractor are defined for x and for y , and the value of z depends on whether the current value of x is closer to the low or high x attractor and on whether the value of y is closer to the low or high y attractor. As in Code 10.14.2, the value of z grows quadratically as x and y deviate from their respective attractors. We use a denser `meshgrid` than before and manually change the view of the graph using the Rotate-3D tool (available when a figure window is active) before copying the figure window and pasting it into a document for presentation outside MATLAB (as in this chapter).

Code 10.14.4:

```

figure(3)
[x,y] = meshgrid(1:61);
[rows columns] = size(x);


x_low_attractor = .5*mean(mean(x));
x_high_attractor = 1.5*mean(mean(x));
y_low_attractor = .5*mean(mean(y));
y_high_attractor = 1.5*mean(mean(y));

k = 5;

for r = 1:rows
    for c = 1:columns
        if abs(x(r,c)-x_low_attractor) <= ...
            abs(x(r,c)-x_high_attractor)
            x_attractor = x_low_attractor;
        else
            x_attractor = x_high_attractor;
        end
        if abs(y(r,c)-y_low_attractor) <= ...
            abs(y(r,c)-y_high_attractor)
            y_attractor = y_low_attractor;
        else
            y_attractor = y_high_attractor;
        end
        z(r,c) = (k*(x(r,c)-x_attractor).^2) + ...
            (k*(y(r,c)-y_attractor).^2);
    end
end

surfc(x,y,z)


```

Output 10.14.4:

The contour map for the surface was partially visible in Output 10.14.2 because we used the `surfc` command rather than the `surf` command. The code for the contour map on its own follows.

Code 10.14.5:

```
figure(4)
contour(x,y,z)
```

Output 10.14.5:

10.15 Generating Rectangular Solids

While considering graphing in three dimensions, it is useful to consider three-dimensional shapes such as rectangular solids, spheres, and cylinders. We can generate rectangular solids in MATLAB using the function shown in Code 10.15.1. A call to that function is shown in Code 10.15.2. The new MATLAB-provided function introduced in Code 10.15.1 is `patch`, which does what `fill` does, but in three as well as two dimensions. Note the specific handles or properties referred to in the `patch` command below. For more information about these and other relevant properties, type `help patch`.

Code 10.15.1:

```
function drawcube=cube(coord);

% coord = 1x3 front/bottom/left coordinates matrix

x = coord(1);
y = coord(2);
z = coord(3);

vertices_matrix = [[x y z];[x+1 y z];[x+1 y+1 z];[x y+1 z]; ...
                  [x y z+1];[x+1 y z+1];[x+1 y+1 z+1];[x y+1 z+1]];


faces_matrix = [[1 2 6 5];[2 3 7 6];[3 4 8 7];[4 1 5 8];...
               [1 2 3 4];[5 6 7 8]];
```

```
drawcube = patch('Vertices',vertices_matrix,'Faces', ...
    faces_matrix,'FaceColor','g');
```

Code 10.15.2:

```
cube([1 2 3])
```

Output 10.15.2:




The image shown above was obtained by using the manual rotation tool in the figure's menu bar, after running the code.

10.16 Generating Spheres and Cylinders

In Code 10.16.1, we use the `sphere` command to generate spheres at different locations. The “sphere” we draw has 24 sides. We put sphere in quotes because 24 sides is many fewer than the infinite number of sides that a true sphere has. The `axis equal` command prevents the spheres from being stretched in the horizontal or vertical dimension, as could occur if MATLAB set the axis automatically. The view of the graph was chosen after using MATLAB's Rotate 3-D tool (available when a figure window is active) before copying and pasting the graphic into the Word document for this chapter. Note that this is the first time we vary where a three-dimensional graphic is placed.

Code 10.16.1:


```
figure(5)
[x y z] = sphere(24);
hold on
for j = 1:2
    surf(x + j,y + j, z + j);
end
axis equal
grid on
box on
view(21,8)
```

Output 10.16.1:

The next program uses the `cylinder` function. We generate two cylinders, one with 24 sides, the other with 18 sides. We place the cylinders at different locations so that one, colored red, seems to sit inside the other, colored blue.

Code 10.16.2:

```
figure(6)
hold on
AZ = -37.5,;
EL = 30;
view(AZ,EL)
for j = 1:2
    if j == 1
        [x y z] = cylinder(24);
        k = 1;
        s = surf(x + k,y + k, z + k);
        set(s,'facecolor','r');
    else
        k = .75;
        [x y z] = cylinder(18);
        s = surf(x + k,y + k, z + k);
        set(s,'facecolor','b');
    end
end
axis off
```


Output 10.16.2:**10.17 Generating Ellipsoids**

Just as a circle is a special kind of ellipse (one whose two axes are of equal length), a sphere is a special kind of ellipsoid (one whose three axes are of equal length). Recognizing that not all axes must have equal length, we can go on to generate ellipsoids, which are useful for depicting biologically relevant forms.

MATLAB provides an `ellipsoid` function. This function returns three matrices, called `x`, `y`, and `z` in the example below. Each matrix is of size `n_facets + 1` by `n_facets + 1`. When rendered with `surf`, the resulting image is an ellipsoid with centers `xc`, `yc`, and `zc`, and radii `xr`, `yr`, and `zr`. The `axis equal` command is used to show the ellipsoid in its intended, stretched form.

Code 10.17.1:

```
xc = 1; yc = 2; zc = 3;  
xr = 1; yr = 1, zr = 3;  
n_facets = 48;  
[x,y,z]=ellipsoid(xc,yc,zc,xr,yr,zr,n_facets);  
surf(x,y,z);  
axis equal;
```


Output 10.17.1:

Humanoid forms can be created by using the `ellipsoid` function, as shown in the following two examples, both of which were written by students in the MATLAB programming seminar where this book was first developed. Code 10.17.2 was written by Matthew Walsh, and Code 10.17.3 was written by Robrecht van der Wel, both of whom gave permission to have their code and outputs reproduced here. Notice that Matt used two commands that have not been discussed so far in this book: `shading interp` and `light`. Robrecht also used three commands not discussed in this book: `shading flat`, `camzoom`, and `camtarget`. The image created with Robrecht's code graced the cover of the first edition of this book. A few modifications of Robrecht's code let us generate the image appearing on the cover of this second edition of *MATLAB For Behavioral Scientists*.

Code 10.17.2:

```
% Ellipsoid_Man_Matt_Walsh
% March_23_2006

close all
clear all
clc

figure(1)
%thorax
[x y z]=ellipsoid(2,3,7.3,1,1,3);
surf(x,y,z);

%head
hold on
[x y z]=ellipsoid(2,3,10.7,1,1,1);
surf(x,y,z);

%shoulder mass
[x y z]=ellipsoid(2,3,9,1,2,.8);
surf(x,y,z);
```

```
%right arm
[x y z]=ellipsoid(3.2,1.4,9.2,1.8,.5,.5);
surf(x,y,z);

%right forearm
[x y z]=ellipsoid(5.9,1.4,9.2,1.3,.4,.4);
surf(x,y,z);

%left forearm
[x y z]=ellipsoid(3.5,4.5,7.1,1.3,.4,.4);
surf(x,y,z);

%left arm
[x y z]=ellipsoid(2,4.5,8.1,.5,.5,1.3);
surf(x,y,z);

%right thigh
[x y z]=ellipsoid(3.33,4,5.1,1.9,.6,.6);
surf(x,y,z);

%left thigh
[x y z]=ellipsoid(3.33,2,4.7,1.9,.6,.6);
surf(x,y,z);

%bubble butt
[x y z]=ellipsoid(2,3,4.7,.8,1.5,.5);
surf(x,y,z);

%right calf
[x y z]=ellipsoid(4.7,2,3,.5,.5,1.4);
surf(x,y,z);

%left calf
[x y z]=ellipsoid(5,2.5,5.2,.5,1.6,.5);
surf(x,y,z);

%left foot
[x y z]=ellipsoid(5.4,1,5.2,1,.2,.505);
surf(x,y,z)

%right foot
[x y z]=ellipsoid(5.2,2,1.8,1,.505,.2);
surf(x,y,z);


grid on
axis on
zlim =[0 20];
```

```

shading interp;
light;
axis equal
set (gca,'view',[107,30], 'AmbientLightColor', [1 0 0]);

```

Output 10.17.2:



Code 10.17.3:

```

% Playing_frisbee_Robrecht_Van_Der_Wel.m
% March_23_2006

close all
clear all
clc

figure(1)
set(gcf, 'Color', [.2 .8 .8]);
title('Playing frisbee', 'FontSize', 20);
colormap(autumn);

subplot(4,2,[1:6]);
% Frisbee person
% Order is: Head, mouth/hair, eyes, nose, shoulders,
% torso, gluteus, left arm, left
% forearm, left hand, right arm, right forearm,
% right hand, right calf, right foot

hold on %Head M/H Eyes Nose Shou Tors GM LA LFA LH RA RFA
RH RC RF
x_1 = [-10 -9.5 -9.2 -9.2 -10 -10 -10 -10 -10 -10 ...
-8.8 -7.3 -7.2 -9.0 -8.5];
y_1 = [3 3.1 3.1 3.1 3 3 3 4.5 4.5 4.5 1.4 ...
2.4 3.9 2.5 2.5];
z_1 = [10.7 10.7 10.7 10.3 9 7.3 4.7 8.1 6.5 5 ...
9.2 9.2 9.2 1.7 .3];

```

```

x_rad_1 = [1 .2 .2 .4 1 1 .8 .5 .4 .3 1.8 ...
.4 .35 .45 .9];
y_rad_1 = [1 .5 1 .2 2 1 .9 .5 .4 .2 .5 ...
1.3 .4 .4 .3];
z_rad_1 = [1 1 .2 .2 .8 3 .5 1.3 1.3 .5 .5 ...
.4 .3 1.5 .2];

for i = 1:length(x_1)
    [xpos_1 ypos_1 zpos_1]= ...
    ellipsoid(x_1(i),y_1(i),z_1(i),x_rad_1(i), ...
    y_rad_1(i),z_rad_1(i));
    surf(xpos_1,ypos_1,zpos_1);
end
shading interp;
light;

[xpos_1 ypos_1 zpos_1]=ellipsoid(-13.1,3.6,1.6,.9,.3,.2);
left_foot_1 = surf(xpos_1,ypos_1,zpos_1);
zdir = [0 1 0];
center = [-13.1,3.6,1.6];
rotate(left_foot_1,zdir,50,center);

[xpos_1 ypos_1 zpos_1]=ellipsoid(-10.5,3.6,3.75,.6,.5,1.3);
left_thigh_1 = surf(xpos_1,ypos_1,zpos_1);
zdir = [0 1 0];
center = [-10.5,3.6,3.75];
rotate(left_thigh_1,zdir,50,center);

[xpos_1 ypos_1 zpos_1]=ellipsoid(-12.1,3.6,2.5,.45,.4,1.5);
left_calf_1 = surf(xpos_1,ypos_1,zpos_1);
zdir = [0 1 0];
center = [-12.1,3.6,2.5];
rotate(left_calf_1,zdir,70,center);

[xpos_1 ypos_1 zpos_1]=ellipsoid(-9.4,2.6,3.8,.6,.5,1.3);
right_thigh_1 = surf(xpos_1,ypos_1,zpos_1);
zdir = [0 1 0];
center = [-9.4,2.6,3.8];
rotate(right_thigh_1,zdir,160,center);

% Catching person
% Order is: Head, hat,mouth/hair, eyes, nose, shoulders,
% torso, gluteus, left arm,
% left forearm, left hand, right arm, right forearm, right hand,
% right thigh,right calf, right foot

```

```
hold on
x_2 = [12 12 11.5 11.3 11.3 12 12 12 12 12 12 11 9.8 8.5 12 ...
      12 11.4];
y_2 = [3 3 3.1 3.1 3.1 3 3 3 1.5 1.5 1.5 4.5 4.5 4.5 ...
      3.5 3.5 3.5];
z_2 = [10.7 11.5 10.5 10.7 10.3 9 7.3 4.7 8.1 6.5 5 ...
      9 9 9 3.9 1.5 .1];
x_rad_2 = [1 1 .2 .2 .4 1 1 .8 .5 .4 .3 1.3 1.3 .5 .6 ...
      .45 .9];
y_rad_2 = [1 1 .5 1 .2 2 1 .9 .5 .4 .2 .5 .4 .2 .5 .4 .3];
z_rad_2 = [1 .2 1 .2 .2 .8 3 .5 1.3 1.3 .5 .5 .4 .3...
      1.3 1.5 .2];

for i = 1:length(x_2)
    [xpos_2 ypos_2
     zpos_2]=ellipsoid(x_2(i),y_2(i),z_2(i),x_rad_2(i),...
        y_rad_2(i),z_rad_2(i));
    surf(xpos_2,ypos_2,zpos_2);
end

[xpos_2 ypos_2 zpos_2]=ellipsoid(11.3,2.4,4,1.3,.5,.6);
left_thigh_2 = surf(xpos_2,ypos_2,zpos_2);
zdir = [0 1 0];
center = [11.6 2.3 2.4];
rotate(left_thigh_2,zdir,55,center)

[xpos_2 ypos_2 zpos_2]=ellipsoid(14,2.5,2.5,.45,.4,1.5);
left_calf_2 = surf(xpos_2,ypos_2,zpos_2);
zdir = [0 1 0];
center = [14 2.5 2.5];
rotate(left_calf_2,zdir,125,center)

[xpos_2 ypos_2 zpos_2]=ellipsoid(14.68,2.5,1.2,.9,.3,.2);
left_foot_2 = surf(xpos_2,ypos_2,zpos_2);
zdir = [0 1 0];
center = [14.68 2.5 1.2];
rotate(left_foot_2,zdir,125,center)

% Playground
[x y z]=cylinder(20,50,1);
surf(x,y,z);
shading flat;

% Frisbee
[x y z] = ellipsoid(0,1,9,1.4,1.4,.2);
surf(x,y,z);
shading flat;
```


```

grid off
axis off
xlabel('x');
ylabel('y');
zlabel('z');

axis equal
set(gca,'view',[134,14], 'AmbientLightColor', [.5,.8,.1]);
camzoom(3);
camtarget([0 0 4]);

```

Output 10.17.3:




10.18 Practicing Plots:

Try your hand at the following exercises, using only the methods introduced so far in this book or in information given in the problems themselves.

Problem 10.18.1:


The previous chapter introduced the `errorbar` function to plot a vertical line relative to points to show the variability of the numbers corresponding to those points. Sometimes behavioral scientists plot one dependent variable against another and both sets of dependent variables have some variability. Write a program that lets you show variability in x as well as in y , similar to the example below. The dummy data used to generate this graph happen to have the property that variability in x and variability in y both scale with their respective means, but that is just an incidental feature of the dummy data.

Output 10.18.1:




Problem 10.18.2:

Adapt the last program to show ellipses around data points. The two axes of the ellipses should correspond to variability along the x and y axes, and the output should resemble the graph below. This problem may take a little detective work on your part if you don't happen to remember the equation for an ellipse. Consult Wikipedia or some other source to find the form of the equation that lends itself most easily to MATLAB coding. The `fill` command was used to generate the white ellipses shown below, which are based on the same data as in Problem 10.18.1.

**Problem 10.18.3:**

The Ebbinghaus illusion is a visual illusion in which two circles of the same size (the two gray circles below) are seen to be of different size depending on the circles around them. Write a program to generate images like those below.

**Problem 10.18.4:**

Adapt your “Ebbinghaus illusion” program so that, from trial to trial, circles of constant size are shown in the central position, and circles of different sizes and positions are shown around the central circles. Write your adapted program so the participant can click on

whichever central circle seems larger. The participant must choose one, so this is an example of a forced choice procedure. Determine the range of outer circle sizes and the range of outer circle distances from the center of the central circle that lead the participant to judge the left central circle as being larger than the right central circle between 25% and 75% of the time.

Problem 10.18.5:

Use `bar3` to visualize the effects of different parameter values on one or more statistical distributions of interest to you (or of your professor). For example the Weibull distribution relates to failure rates over time and has been applied to such things as the characterization of infant mortality rates. Wikipedia or other sources can be used to obtain information about statistical distributions. For example, Wikipedia includes the following statement in its (August 27, 2006) article about the Weibull distribution: “Given a random variate U drawn from the uniform distribution in the interval $(0, 1)$, then the variate

$$X = \lambda(-\ln(U))^{1/k}$$

has a Weibull distribution with parameters k and λ . This follows from the form of the cumulative distribution function.” Show the effect of k and λ on X in a three-dimensional bar graph.

Problem 10.18.6:

Draw on the code in Sections 10.9–10.14 to generate one or more 3-D graphs that show real or simulated data for a behavioral science problem of interest to you (or your professor).

Problem 10.18.7:

Draw on the code in Section 10.15 to depict a staircase with a railing.

Problem 10.18.8:

Draw on the code in Section 10.17 to show a humanoid descending the staircase, or in some other pose that might be useful to you in your research.

Problem 10.18.9:

Repeat the demonstration of Section 10.4, but change the proportion of pixels that are white and black. How does this affect the appearance of the squares? Add multiple internal squares at different apparent depths and/or vary the contrast between the pixels by modifying the contents of the color map.

Problem 10.18.10:

MATLAB does not have a good way to make patterned bar graphs that show up well in grayscale print. Using what you know about lines, explore how to superimpose a pattern of diagonal lines on one of the bars in such graph.

Problem 10.18.11:

Generate the rainbow of Output 10.4.1 using a $I \times 6$ image array and a 6×3 color map that defines the colors for each of the six cells.

11. Animation and Sound

This chapter covers the following topics:

- 11.1 Animating by changing successive images
- 11.2 Watching comets
- 11.3 Animating by drawing now
- 11.4 Making movies
- 11.5 Saving movies
- 11.6 Reading and running previously saved movies
- 11.7 Playing beeps
- 11.8 Loading and playing sound files
- 11.9 Controlling volume
- 11.10 Staggering or overlapping sounds and delaying sounds
- 11.11 Controlling volume while staggering or overlapping sounds
- 11.12 Creating your own sound files computationally
- 11.13 Writing and reading files for sound
- 11.14 Learning more about sound-related functions
- 11.15 Practicing animation and sounds

The commands that are introduced and the sections in which they are premiered are:

<code>delete</code> (<i>handle</i>)	(11.1)
<code>comet</code>	(11.2)
<code>comet3</code>	(11.2)
<code>drawnow</code>	(11.3)
<code>getframe</code>	(11.4)
<code>movie</code>	(11.4)
<code>movie2avi</code>	(11.5)
<code>VideoReader</code>	(11.6)
<code>VideoWriter</code>	(11.6)
<code>beep</code>	(11.7)
<code>sound</code>	(11.8)
<code>soundsc</code>	(11.9)
<code>audioplayer</code>	(11.10)
<code>play</code>	(11.10)

playblocking	(11.10)
timer	(11.10)
wavread	(11.11)
wavwrite	(11.11)
audioread	(11.13)
audiowrite	(11.13)
audiorecorder	(11.14)

11.1 Animating by Changing Successive Images

Seeing things change can help you understand them better and can also help you appreciate them more from an aesthetic standpoint. In this chapter, we build on this observation by delving into animation. First we apply what we have covered about graphics to create moving images. Then we turn to some tools that MATLAB provides for creating, reading, and saving animations in ways that afford professional-looking dynamic renderings. The last parts of the chapter concern sounds.

The essence of computer animation, like the essence of traditional cinema, is the display of series of images (“frames”) presented at sufficiently short intervals to be perceived as moving (or holding still if the apparent positions of the depicted objects remain the same). To use MATLAB for computer animation, you can take advantage of the way MATLAB represents the component parts of a graph or figure. As each component of a graph or figure is drawn, you can optionally remember the value of its handle variable, and you can then use that handle variable to change the characteristics of the object on the screen, including, if you wish, removing the object by deleting its handle.

Here is an example of a program that puts a marker on the screen, then erases it, and then draws another marker, slightly larger, at regular intervals of .1 seconds. The marker is moved to the right by a small amount each time it is redrawn. Successive plots to the same window replace the prior plot. The `hold off` command reinforces this idea but is not strictly required. The animation cannot be shown in this book, of course, but you can see it if you run the program.

Code 11.1.1:

```
figure(1)
clf
for loopvalue = 2:2:38
    thisx = 1 + loopvalue/10;
    thisy = 2 - loopvalue/20;
    thissize = loopvalue;
    plot(thisx,thisy,'*','Markersize', thissize);
    hold off
    axis([1 5 0 2])
    pause(.1);
end
```

The next program creates a series of images of a very simple “arm” moving from one position to another, with sufficiently short inter-image delays to give the illusion of motion, using a slightly different method, namely, repeatedly plotting and erasing the figure. The arm’s “shoulder” is always located at position (ShoulderX, ShoulderY), and the shoulder’s angle, ShoulderAngle(i), moves in six steps of equal size from $.15\pi$ to $.45\pi$. The arm’s “elbow” is located, at each moment, i, at position (ElbowX(i), ElbowY(i)), depending on ShoulderAngle(i). Similarly, the arm’s “hand” is located, at each moment, i, at position (HandX(i), HandY(i)), depending both on ShoulderAngle(i) and ElbowAngle(i), which moves in six steps of equal size from $.75\pi$ to $.55\pi$. We plot the y values of the shoulder, elbow, and hand against the x values of the shoulder, elbow, and hand for each move. To keep the axes the same in successive plots, we use hold on and we set xlim and ylim to visually satisfying values. To ensure that we can see the figure as the animation unfolds, we use the pause command before the first plot command is issued, remembering to look at the figure window while hitting whichever key we choose to terminate the pause. We first plot the x, y data with black circles and lines ('ko-'), then pause for .2 seconds, and then erase the line by deleting its handle. Recall from Chapter 9 that any plotted object can be assigned a handle. Given this capability, we use the delete command to remove (erase) the plotted object from the figure without affecting any other objects in the plot. Pausing for .2 seconds and then deleting the just-plotted circles and lines only occurs if want_animation is true. If want_animation is false, a set of superimposed plots is created that can be copied and reproduced elsewhere—for example, as a figure to be published in a paper or in this book (Output 11.1.2).

Code 11.1.2:

```
close all
clear all
shg

ShoulderX = 0;
ShoulderY = 0;

moves = 6;
want_animation = true;

ShoulderAngle = linspace(.15*pi, .45*pi, moves);
ElbowAngle = linspace(.75*pi, .55*pi, moves);

position = [];


figure(1)
hold on; grid on; box on;
xlim([-2.5 2.5]);
ylim([-2.5 2.5]);

for i = 1:moves
    ElbowX(i) = ShoulderX + cos(ShoulderAngle(i));
    ElbowY(i) = ShoulderY + sin(ShoulderAngle(i));
```

```
HandX(i) = ElbowX(i) + cos(ElbowAngle(i));
HandY(i) = ElbowY(i) + sin(ElbowAngle(i));

position = [position; [ShoulderX ElbowX(i) HandX(i)] ...
              [ShoulderY ElbowY(i) HandY(i)]];
armhandle = plot(position(i,1:3),position(i,4:6),'ko-');
if want_animation
    pause(.2)
    if i < moves
        delete(armhandle);
    end
end
end
if not(want_animation)
    saveas(gcf,'Output_11_1_2','eps')
end
```

Output 11.1.2:



11.2 Watching Comets

MATLAB provides other, more automatic, ways of creating dynamic motion. One is to use MATLAB's `comet` command, which displays a moving object along with a trailing tail as it streaks across a plane. MATLAB's `comet3` command displays an object moving, or seeming to move, in three dimensions rather than two.

It is impossible to do justice to the animations that can be achieved with `comet` and `comet3` in the pages of this book. We encourage you to read about these commands in `doc comet` and `doc comet3`. You can copy the code from there and run `comet` and `comet3` to admire the resulting “heavenly” animations.

11.3 Animating by Drawing Now

Ordinarily, when you plot a number of things on a graph, the figure image is not updated until there is a “break in the action,” such as a `pause` statement. While this is an efficient

strategy most of the time—you don’t have to wait through many screen updates to get to the end of the plot—for some animations you may want to see the result of each plot more immediately. Here the `drawnow` command is useful. It forces immediate rendering of the new image. When `drawnow` is contained in a `for` loop with more than one cycle, the immediate rendering occurs each time through the loop. Furthermore, with `hold off`, what was rendered before is not retained. When Code 11.3.1 is run, figure (2) appears and the arm is seen to move from its starting position to its final position, where it remains until a key is hit again to terminate the final pause. If `drawnow` were omitted, no action would be seen, since the figure would be updated on the screen only at the end of the plotting. We use a pause after each frame to slow the movie down.

Code 11.3.1:

```
figure(2)
shg
hold off
for i = 1:moves
    plot(position(i,1:3),position(i,4:6),'ko-');
    grid on
    xlim([-2.5 2.5]);
    ylim([-2.5 2.5]);
    drawnow
    pause(.3)
end
close(2)
```

11.4 Making Movies

If you generate animations with MATLAB, it’s nice to share them with others, even those who don’t necessarily use MATLAB themselves. Is there a way to save an animation as a movie that can viewed outside MATLAB, say in Windows Media Player or QuickTime?

There is a way to do this, as shown in Code 11.4.1. As you may guess, a movie is generated by first generating each of the pictorial frames that constitute the movie, and then putting those frames together, in sequential order, in a file that can be played as a movie. This code uses three new features. One is a parameter of `plot` called `erasemode`, which is set to `normal` to ensure that the plot is displayed as wished in this context. The second is the command `getframe`, which assigns the contents of the current figure window to the current frame. The third is `movie`, which displays the frames obtained through `getframe`. Notice that the call to `movie` in Code 11.4.1 has two arguments. The first, which is obligatory, is the variable that contains the frames to be shown—in this case, `F`. The second argument, which is optional, is the number of times the movie will be shown.

A peculiar feature of `movie` is that the frames being loaded into the movie are shown while the loading occurs. Thus, making 1 the second argument of `movie` shows the movie *twice*, once (slowly) while it is being generated and then again (at full speed) while it is being “officially shown.”

Code 11.4.1:

```
grid on
box on
hold on
for i = 1:moves
    plot(position(i,1:3),position(i,4:6),...
         'ko-', 'erasemode', 'normal');
xlim([-2.5 2.5]);
ylim([-2.5 2.5]);
    F(i) = getframe;
end
pause (1)
movie(F,1)
```

11.5 Saving Movies

Having made a movie, you will need to save it if you want to retrieve it later. A single command achieves this: `movie2avi`. As shown in Code 11.5.1, `movie2avi` has two arguments. The first is the name of the file being saved. The second is the name of the target file. The to-be-saved file name is a string and should have the `.avi` suffix.

Code 11.5.1:

```
| movie2avi(F, 'ArmMove.avi')
```

Once this code has been run, you can confirm that the file can be opened and viewed outside of MATLAB (e.g., in Windows Media Player or in PowerPoint). Note that the `.avi` format is a Windows-specific format. Macintosh users may have to install an appropriate media player (e.g., WMV or Flip Player) to view the movies outside of MATLAB, even if they were generated on the Mac.

A recent addition to MATLAB (as of release 2013a), `VideoWriter` allows `.avi` and other movie formats to be generated and run on all platforms. Code 11.5.2 uses the frames just generated to make the movie run in reverse.

Code 11.5.2:

```
% Write the movie backwards
figure(3);
clf
writerObj = VideoWriter('evoMmrA'); % Release 2013a or later
open(writerObj)
for k = 6:-1:1
    image(F(k).cdata)
    frame = getframe;
    writeVideo(writerObj, frame);
end
close(writerObj) % 'evommrA.avi' is readable in Mac OS, too
```

11.6 Reading and Running Previously Saved Movies

Much as it is desirable to save .avi files for later use, it is desirable to be able to read and run previously saved movies, including ones not generated in MATLAB. In this section we first read and run the shuttle.avi file (distributed with MATLAB release 2013a and later) using VideoReader. We use VideoReader to read the file and determine the number and dimensions of the images (frames) it contains, which we find in the structure (see Chapter 7) that we call myMovieObj. Then, we create a new array of frames, myFrames, that we can examine individually or use for other purposes. In Code 11.6.1 we examine every tenth frame of shuttle.avi.

Code 11.6.1:

```
myMovieObj = VideoReader('shuttle.avi');
nFrames = myMovieObj.NumberOfFrames;
for k = 1 : nFrames
    myFrames(k).cdata = read(myMovieObj, k);
end
for k = 1:10:nFrames
    image(myFrames(k)); shg;
    text(50,50,sprintf('Frame Number: %d',k));
    pause(0.5);
end
```

You could perform the same operation on any .avi (or .mpg or .mov) file obtained on the web or from a colleague, even if it had not originally been generated through MATLAB. Code 11.6.2 reads the video, shuttle.avi, and makes a new video in which the order of frames is reversed. The video data are in a $x \times y \times 3 \times f$ matrix, where x and y (the first two dimensions) are the width and height of the image, the third dimension is of size 3 to represent the red, green and blue value of each pixel, and f (the fourth dimension) is the number of frames in the original.

Code 11.6.2:

```
xyloObj = VideoReader('shuttle.avi');
vidFrames = read(xyloObj);
nFrames = size(vidFrames,4);
figure(4)
writerObj = VideoWriter('elttuhs.avi');
open(writerObj);
for k = nFrames:-1:1
    image(vidFrames(:,:,k))
    frame = getframe;
    writeVideo(writerObj,frame);
end
close(writerObj)
```


You can also examine any single frame (the sixth, say) of `shuttle.avi`, taking advantage of the fact that the frames are now in the array, `vidFrames`. If you assign `vidFrames(:, :, :, 6)` to another variable, `im`, the subsequent commands `image(im)` and `axis image` will show that frame.

Code 11.6.3:

```
im = vidFrames(:, :, :, 6);  
figure(2)  
image(im)  
axis image
```

11.7 Playing Beeps

We now direct your attention to a modality that has gotten scant coverage in this book, though that modality has undoubtedly captured your attention on many occasions if you have written code that happened to have problems. We refer to sound, and specifically to the beeps you have probably heard alerting you to errors picked up by the MATLAB compiler. It would be nice to be able generate sounds other than, or in addition to, beeps, and also to do so through means other than erring.

Our first sound-generation program generates two beeps. The `beep` command is given, there is a `pause` for 2 seconds, and then the `beep` command is re-issued.

Code 11.7.1:


```
beep  
pause(2)  
beep;
```

11.8 Loading and Playing Sound Files

The second sound example shows how sounds can be generated using files that come with MATLAB: `chirp.mat`, `handel.mat`, and `gong.mat`. To hear the chirping, load the `chirp.mat` file and then issue the `sound` command. In so doing, you will take advantage of the fact that when `chirp` is loaded, the variables `y` and `Fs` are automatically assigned. To see what the `chirp` data look like in graphical form, you can `plot` the sound data, `y`. To see how the data of `chirp.mat` are internally represented, you can use `whos` to reveal the properties of `y` and `Fs`.

Code 11.8.1:

```
load chirp  
sound(y)  
plot(y, 'k')  
commandwindow  
whos  
Fs
```

Output 11.8.1:**Output 11.8.2:**

Name	Size	Bytes	Class	Attributes
Fs	1x1	8	double	
y	13129x1	105032	double	

Fs =
8192

From Output 11.8.2, you can see that the values in the chirp file are double precision real numbers (see Chapter 7) within the range -1 to +1 occupying a matrix of 13,129 rows and 1 column. The sampling rate of the sound is in Fs, in samples/second. The duration of chirp-ing is thus 13129/8192, or around 1.5 seconds.

11.9 Controlling Volume

The third code example shows how to control the volume of a played sound file. Here we load the sound data file called `handel.mat`. Knowing that the output of `load handel` is `y`, we supply `y` as the first argument to a function called `soundsc`, which stands for “sound, scaled.” The second argument of this function is a matrix whose minimum and maximum values determine the volume of the generated sound. The closer these minimum and maximum values are to zero, the greater the volume. (Yes, that last statement is correct, though it is counter-intuitive.) Meanwhile, we pause 9 seconds, giving Handel’s Hallelujah chorus (at least this short excerpt) time to finish before playing it again more softly with more extreme values for the second argument of `soundsc`. The excerpt’s duration (8.92 seconds) was computed by dividing the number of samples, 73,113, by the sampling rate, 8192/second, both of which we determined by examining the variables in the Command window after `handel.mat` was loaded.

Code 11.9.1:

```
load handel
soundsc(y, [-3.25 3.25])
pause(9)
soundsc(y, [-15.25 15.25])
```

Listening to the output indicates that the sound file is played at different volumes depending on the second argument of `soundsc`. Regardless of the volume assigned in `soundsc`, the original value of `y` is unchanged, so `soundsc`, at least as used here, only serves as an “external volume controller.”

11.10 Staggering or Overlapping Sounds and Delaying Sounds

In Code 11.9.1, we delayed the second presentation of the excerpt of the Hallelujah chorus by pausing for 9 seconds, so the first presentation could finish. MATLAB provides two other sound functions, called `play` and `playblocking`, which let you control the staggering or overlapping of sounds more directly. `play` and `playblocking` each take three arguments. The first is a variable representing the sound to be played. The next two arguments are optional and indicate the beginning and ending samples to play.

We begin by setting up two `audioplayer` objects, one for `handel` and one for `chirp`. `audioplayer` objects are structures that hold all of the relevant data about a sound object. Then, we play the two samples sequentially, using `play`, the command for playing `audioplayer` objects. After we start the `handel` sound, we pause for only 2 seconds, then start the `chirp` sound. Notice that the chirping birds “join the chorus” and overlay the `handel` sound after the 2 seconds have elapsed.

Code 11.10.1:

```
load handel;
handelplayer = audioplayer(y,Fs);

load chirp;
chirpplayer = audioplayer(y,Fs);

play(handelplayer)
pause(2)
play(chirpplayer)
```

What if we did not want to go on to the chirps until the chorale ended? In that case we would use `playblocking`. This function waits until the entire selection is finished before going on. The result is all of the `handel` sample, the pause of 2 seconds, and then all of the `chirp` sample.

Code 11.10.2:

```
playblocking(handelplayer)
pause(2)
playblocking(chirpplayer)
```

It is important to appreciate that the differing effects of `play` and `playblocking` don’t only apply to the staggering or overlapping of *sounds*; they also apply to the staggering or overlapping of other events. Thus, if you want to plot points, display images, or read in keystrokes using `ginput` (see Section 10.5) while sounds are being played, you can use `play`. If you prefer to wait, use `playblocking`.

You can also control a delay before a sound begins to play with the `timer` function, which can be used to specify that a particular action be initiated in the future. Here `play(handelplayer)` will be executed 1.5 seconds after the timer starts.

Code 11.10.3:

```
t = timer('TimerFcn','play(handelplayer)', ...
        'StartDelay', 1.5);
start(t)
```

11.11 Controlling Volume While Staggering or Overlapping Sounds

Section 11.9 showed how to control volume with the `sound` command. However, the `sound` command doesn't let you easily control the staggering or overlapping of sounds. On the other hand, `play` and `playblocking` let you easily control the synchrony or asynchrony of sound files. This raises the question of whether there is a way to control the volume while using these commands, so you have the best of both worlds—a command that lets you control the synchrony or asynchrony of sounds as well as their volumes. A solution follows.

We use `playblocking` so each sound begins only when the prior one has finished, and we scale the variable, `y`, loaded from `gong.mat` by different amounts in three `audioplayer` objects. If you run this program on your computer, you will hear a loud gong, a soft gong, and then a medium-amplitude gong. You can use this example as a basis for controlling the volume of other sound files.

Code 11.11.1:

```
load gong; % Loads y and Fs for sound from gong.mat
tooloudplayer = audioplayer(y,Fs); % Volume is too loud!
toosoftplayer = audioplayer(y/5,Fs); % Volume is too soft!
goldilocksplayer = audioplayer(y/2,Fs); %Volume is just right!

playblocking(tooloudplayer);
playblocking(toosoftplayer);
playblocking(goldilocksplayer);
```

11.12 Creating Your Own Sound Files Computationally

The graph in Output 11.8.1 is familiar-looking plot of a one-dimensional matrix. Can such data serve as inputs to `sound` or `play`? Can you, in other words, *listen* to your data files as well as *see* them? The answer, you will be happy to hear, is yes.


Code 11.12.1 shows how to generate a data file that serves the somewhat mundane function of creating static. Having participants listen to static is often useful in behavioral research, particularly if you want the participant not to hear other sounds in the environment.

The particular form of static that is generated here is white noise. A white-noise signal is one for which the intensities of all frequencies within the included band of frequencies is the same. You can create a reasonable approximation to such a signal with a uniform distribution, using the `rand` function (see Section 4.4). Using `rand`, you can create a $1 \times n$ matrix called `noise`, where `n` determines the duration, `d`, of the sound you wish to generate—here `d = 1.0` s—multiplied by the sample frequency, `sf`, which is here set to 22050 Hz (samples per second). In the code below, we normalize the values of `noise` so they occupy the range 0 to 1 because we know that the `sound` function works best with values between -1 and $+1$. We issue the `sound` command, which converts the data comprising the `noise` matrix to auditory energy at a sample frequency `sf`. Finally, we plot `noise` over the entire sample interval in the top subplot, and expand it (just the first 250 samples) so we can see the details of the noisy signal in the bottom subplot.

Code 11.12.1:

```
sf = 22050;                % sample frequency
d = 1.0;                  % duration
n = sf*d;                 % number of samples
noise = rand(1,n);        % uniform distribution
noise = noise / max(abs(noise)); % normalize
sound(noise,sf);          % play sound
subplot(2,1,1)
plot(noise,'k')
xlim([1 n]);
subplot(2,1,2)
plot(noise(1:250),'k')
```

Output 11.12.1:




In the next example, we generate a sine wave, both to see and hear it. The structure of the program is similar to the one used to generate static. However, the data file comprising the first argument to `sound` is a sinusoidal function rather than a uniform distribution.

After issuing the `sound` command, we generate two subplots. The top one shows the full sinusoidal function, which, like the noise plot, is so densely packed that it looks like a solid bar. The bottom subplot shows just the first 250 values of `s`, showing more clearly the periodic oscillation characteristic of a sine wave. Listening to the sine wave reminds us that periodic oscillations are called pure tones. In this case, because we set the carrier frequency to 440 Hz, the pure tone we hear is the note A4, or “middle A” on a piano. This is the note to which classical musicians generally tune their instruments.

Code 11.12.2:

```
cf = 440; % carrier frequency (Hz)
sf = 22050; % sample frequency (Hz)
d = 1.0; % duration (s)
n = sf * d; % number of samples
s = (1:n) / sf; % time-dependent values
tone = sin(2 * pi * cf * s); % sinusoidal modulation
sound(tone,sf); % sound presentation
subplot(2,1,1)
plot(tone,'k')
subplot(2,1,2)
plot(tone(1:250), 'k')
```

Output 11.12.2:




In Code 11.12.3, we again generate a sine wave, but this time we let the intensity grow as time passes. We do this by defining a value, `a`, that increases linearly from `1/length(tone)` up to 1, with as many steps as `length(tone)`. Then we show both the overall waveform and the detail from the tone’s first part.

Code 11.12.3:

```

a = linspace(1/length(tone),1,length(tone));
sound(a.*tone,sf)
plot(a.*tone,'k')
subplot(2,1,1)
plot(a.*tone,'k')
xlim([1 n]);
subplot(2,1,2)
plot(a(1:5000).*tone(1:5000),'k')
ylim([-1 1]);

```

Output 11.12.3:

The last example in this section is adapted from a program in the public domain at http://users.ece.gatech.edu/~bonnie/book/OnlineDemos/SignalsAndSounds/synthetic_music.html. The program lets you generate a C major scale by defining the notes in the scale relative to A4. At the end, we save the file for later use with the `audiowrite` command.

Code 11.12.4:

```

fs = 8000;           % sampling frequency
t = 0:1/fs:0.25;    % length of each note
tspace = 1.0;       % length of pause between notes
fr = 2^(1/12);       % frequency ratio between neighboring keys
A4 = 440;            % reference note for others
B4 = A4*fr^2;
C4 = A4*fr^(-9);
D4 = A4*fr^(-7);
E4 = A4*fr^(-5);
F4 = A4*fr^(-4);

```

```

G4 = A4*fr^(-2);
C5 = A4*fr^3;
xspace = zeros(1,tospace*fs);    % set pause
x = [cos(C4*2*pi*t),xspace, ...
     cos(D4*2*pi*t),xspace, ...
     cos(E4*2*pi*t),xspace, ...
     cos(F4*2*pi*t),xspace, ...
     cos(G4*2*pi*t),xspace, ...
     cos(A4*2*pi*t),xspace, ...
     cos(B4*2*pi*t),xspace, ...
     cos(C5*2*pi*t)];
myScale = audioplayer(x,fs);
play(myScale)
audiowrite('scale.wav',x,fs)

```

The foregoing example shows that a cosine function yields tones that are “just as pure” as a sine function. This to be expected because a cosine function is just a phase-shifted version of its corresponding sine function. Another point illustrated by the foregoing example is that `play` and `playblocking`, used in conjunction with `audioplayer`, can be used to play generated files, just as sound can be.

11.13 Writing and Reading Files for Sound

The final matter to be addressed here is how files for sound can be written to external files and in turn be read from such files. At the end of Code 11.12.4 we used the `audiowrite` command to write the sound data (`x`) and sampling frequency (`fs`), to an external file. The name of the external file is a string consisting of the name of the file—in this case `scale`—followed by `.wav`, which identifies the file type. After writing the data to the file using `audiowrite`, we can read the file using `audioread` to read in sound and sampling frequency variables (`y` and `Fs`, respectively). Finally, we play the file that was read in, with the `sound` command.

Code 11.13.1:

```

| [y,Fs] = audioread('scale.wav');
| sound(y,Fs)

```

11.14 Learning More About Sound-Related Functions

As always with MATLAB, there are other methods that can be used in conjunction with topics covered here. To learn how to record sounds using your computer’s microphone, explore `audiorecorder`. All the sounds we have described have been represented by $1 \times n$ matrices. You can also explore how to use $2 \times n$ matrices to play stereophonic sounds or sounds that have entirely different content for the two ears. MATLAB’s Help will give you an earful on that!

11.15 Practicing Animation and Sounds

Try your hand at the following problems, using only the methods introduced so far in this book or in the problems themselves.

Problem 11.15.1:

Write an animation program to show the view from the flight deck of a starship entering warp speed, so that the figure shows an expanding optic flow field of multiple objects, looming closer. Each point will appear to grow in size as it follows a straight trajectory toward the edge of the screen. (Hint: After you have plotted an array of points, use `get` in a `for` loop to change the size and position of each of the points).

Problem 11.15.2:

Adapt the program used to generate the motion of a right arm (Code 11.1.2 and 11.3.1) so the left arm and right arm both move at once. Save the output as a movie so it can be viewed outside MATLAB.

Problem 11.15.3:

Adapt the program used to generate the motion of a right arm so one arm or both arms (as you wish) reach out to contact a moving ball. Save the output so it can be viewed outside MATLAB.

Problem 11.15.4:

Write a program for an experiment on intermodal perception. For example, show an animation along with a sound sequence that either fits or does not fit with the animation. Such stimuli have been presented to infants to determine whether infant gaze durations depend on the match between visual and auditory stimuli.

Problem 11.15.5:

Write a program to read and run a previously saved movie either with the frames in their original order, in reverse order, or in some scrambled order. Save the output so it can be viewed outside MATLAB.

Problem 11.15.6:

Adapt Code 11.12.4 to play a melody such as “Twinkle, Twinkle, Little Star.” Use functions so you can specify the sounds economically (i.e., as a string of note values, such as “CCGGAAG”), and easily change the tune.

Problem 11.15.7:

Write a program to take a series of tones of different notes and durations and play it repeatedly, at a slow tempo, so you can play along as you learn your musical instrument. Adjust the tempo as you master the melody.

Problem 11.15.8:

In working on the last problem, you may have noticed that there is a bit of a “click” between the tones. This is due to the abrupt transition from zero to full amplitude of the waveform. Devise a way for each sound to gradually begin and end, so the onset and offset of each sound are more gradual.

Problem 11.15.9:

Write a program for an experiment in which participants make auditory discriminations. For example, participants perform a forced choice task in which they indicate which of two tones is louder the first or the second.

Problem 11.15.10:

Write a program in which subjects answer questions and get auditory feedback that indicates whether they got the answer right or wrong. To make things a bit fancy, change the volume of the sound according to how quickly the question was answered and according to whether the answer was correct or incorrect.

12. Enhanced User Interaction

This chapter covers the following topics:

- 12.1 Getting less “clunky”
- 12.2 Creating graphic user interfaces (GUIs)
- 12.3 Using built-in user interface utilities
- 12.4 Writing code for user interface functions
- 12.5 Prototyping user interfaces using GUIDE
- 12.6 Recording user interactions
- 12.7 Practicing interfaces and interactions

The commands that are introduced and the sections in which they are premiered are:

<code>errordlg</code>	(12.2)
<code>inputdlg</code>	(12.2)
<code>listdlg</code>	(12.2)
<code>msgbox</code>	(12.2)
<code>questdlg</code>	(12.2)
<code>uigetdir</code>	(12.2)
<code>uigetfile</code>	(12.2)
<code>uiopen</code>	(12.2)
<code>uicontrol</code>	(12.3)
<code>@</code>	(12.4)

12.1 Getting Less “Clunky”

So far in this book, all the interfaces and interactions with users have been a bit “clunky.” For example, in Section 6.2, you saw how to use `input` to collect keyboard responses from users (participants in experiments or surveys), but the interface was quite bare bones. The user effectively sat at the keyboard, much as you do when you program. In Section 10.5, you saw how to use `ginput` to record where users click in figure windows. The resulting interface may have been a bit less ascetic than the MATLAB command line, but it still lacked the bells and whistles, or at least the aesthetic feel, you expect when you interact with computers and other computer-driven devices.

Computer programs that show text and graphics that allow for clicking on buttons, scrolling up and down, typing in numbers, and so on, provide you with a Graphical User Interface or GUI (pronounced “goeey”). This chapter acquaints you with MATLAB-based GUIs (see Sections 12.1–12.4), including MATLAB’s GUIDE tool for constructing user interfaces. The chapter then turns back to the non-GUI world (where we three authors have mostly worked), showing some ways to have more flexibility and power in the kinds of user interactions your programs can support (see Sections 12.5–12.6).

12.2 Creating Graphic User Interfaces (GUIs)

MATLAB lets you set up GUIs in three ways. First, it affords a number of built-in interface functions that provide many of the interface needs you may have. These, combined with the graphical capabilities described in Chapter 9, enable interaction between the user and the program. The built-in functions take care of a number of commonly encountered features of GUI programming, allowing you, the programmer, to work at a high level. (Recall the discussion of the advantages of working at a high level in the introduction to Chapter 8 on Modules and Functions.)

Second, MATLAB lets you work at a detailed level, using the `uicontrol` function. Via `uicontrol` you can program details of interface functions that you might want. Such details include aspects of GUI size, location, and text.

Third and finally, MATLAB provides a “drag and drop” utility for constructing GUIs. This utility, called GUIDE, lets you place user controls in windows where you want them. Then it automatically generates MATLAB code corresponding to their placement and control.

For behavioral scientists, there is another important reason to understand how user interface interactions are implemented in MATLAB. In the laboratory, you often need to monitor participants’ performance by recording their response selections, reaction times, and other dependent measures. Many of these needs can be addressed by “going GUI.”

12.3 Using Built-In User Interface Utilities

To help you make your way into GUI programming, it will help you to know what built-in utilities are available. These built-in utilities can be found by opening MATLAB’s documentation and searching for “predefined dialog boxes.” Here are some of the dialogs that appear when you do this:

<code>errordlg</code>	Create and open error dialog box
<code>inputdlg</code>	Create and open input dialog box
<code>listdlg</code>	Create and open list-selection dialog box
<code>msgbox</code>	Create and open message box
<code>questdlg</code>	Create and open question dialog box
<code>uigetdir</code>	Open standard dialog box for selecting directory
<code>uigetfile</code>	Open standard dialog box for retrieving files
<code>uiopen</code>	Interactively select file to open and load data
<code>uiputfile</code>	Open standard dialog box for saving files
<code>uisave</code>	Interactively save workspace variables to MAT-file

These functions can add a great deal of flexibility to your program development. They make it easy to get information from users, such as which files to use and where to save the outputs. The interface utilities make it possible to control and constrain the responses that users give, so only valid responses can be selected. The utilities also help remind users of what information is needed from them. These interactions occur in dialog boxes. We can’t show their operation in print, but you can get a feel for them by typing commands in the Command window and observing what happens

when the commands are executed. For example, type `msgbox('Let this be a warning!', 'Warning Message', 'warn')` in the Command window. The first argument is a string with the message content. The second is the name of the dialog box. The third specifies the type of icon for the message. In this case, `warn` means the message is a warning alert. `help` or `error` in this argument would produce a “Help” or “Error” icon, respectively.

Code 12.3.1 demonstrates the built-in functions `questdlg` and `msgbox`. Suppose your program is going to save some data in a file named `File1.txt`. Because an older file of this name might already exist the program should first use `exist('File1.txt', 'file')` to check whether there already is a file or folder with that name in the present folder. (The argument `file` restricts the test to files or folders.) A returned value of 2 indicates that such a file exists, in which case the program then uses `questdlg` to ask if it is OK to overwrite the old file, and it follows up with `msgbox` to reassure the user that the old file won't be deleted if the user elects to overwrite.

Code 12.3.1:

```
% Construct a File1.txt to use for the example
fileout1 = fopen('File1.txt','wt');
fprintf(fileout1,'THIS IS THE CONTENTS OF THE FIRST FILE\n');
fclose(fileout1);

% Now test for the existence of File1.txt
if exist('File1.txt','file') == 2
    mybutton = questdlg('Delete old File1.txt? ','File1','Yes');
    switch mybutton
        case 'Yes'
            delete('File1.txt');
            msgbox('File1.txt deleted!')
        case {'No' 'Cancel'}
            msgbox('File1.txt unchanged ... Exiting','', 'warn')
            return
    end
end
end
```

Code 12.3.2 opens an existing file, modifies it, then saves it under a new name. Assuming you don't know ahead of time which file the user wishes to modify, the program uses `uigetfile` to select a file, using the standard File Open dialog box, specifying that files of the type `.txt` are the default for selection. After the user selects `File1.txt`, it then reads the text from `File1.txt`, modifies its contents (one string), and writes the modified string to `File2.txt`. Finally, it confirms that the operation has been completed with a call to `msgbox`. These commands are presented below. We have omitted the semi-colon from the call to `uigetfile` to demonstrate the value returned from it (selecting `file1.txt` in the `uiopen` dialog box).

Code 12.3.2:

```
% Make sure File1.txt exists for the example
fileout1 = fopen('File1.txt','wt');
fprintf(fileout1,'THIS IS THE CONTENTS OF THE FIRST FILE\n');
fclose(fileout1);

% Read in a file, change it, and write out the modified file.
infilename = uigetfile('*.txt')
theText = fileread(infilename);
theText = strrep(theText,'FIRST','SECOND');
fileout2 = fopen('File2.txt','wt');
fprintf(fileout2,theText);
fclose(fileout2);

% Verify that the updated text is now in File2.txt
msgbox('File2.txt successfully created','All done!','Help')
fprintf('\nContents of File2.txt:')
type File2.txt
```

Output 12.3.2:

```
infilename =
File1.txt

Contents of File2.txt:
THIS IS THE CONTENTS OF THE SECOND FILE
```

12.4 Writing Code for User Interface Functions

The dialog boxes just discussed appear sequentially on the screen at fixed locations. Their order of appearance is predetermined by the program. By contrast, many familiar applications have a window with several controls active at once. These controls can include pop-up menus and buttons that can be used by user in any order he or she prefers. How can you set up such an interface?

You can do so with the `uicontrol` function. This function implements many graphic interface elements you are familiar with:

checkbox	A square box that can be checked or unchecked
edit	A text field that can be edited
listbox	A menu from which one or more items can be selected
popupmenu	A pop-up menu from which one item can be selected
pushbutton	A button on which the user can click
radiobutton	A group of buttons of which one and only one can be selected

slider	A control that can be moved to indicate a value
text	A text field that is fixed (cannot be edited)
togglebutton	A button that alternates its state when pressed


When these kinds of controls are used, the program has special *callback* functions that are never called by the main function. Rather, each callback function is idle until it is directly activated by one of the user's actions (such as a click) in the program window.

An example of `uicontrol` follows. Here the button in the window executes a beep when the button is pressed. The program runs for 10 seconds, beeping for every button press, and then quits, closing the figure.

Code 12.4.1:

```
h = figure;
set(h, 'position', [ 427  306  512  100])
hpb = uicontrol('Style', 'pushbutton', ...
    'String', 'Make a sound!', ...
    'Position', [20 20 150 20], ...
    'Callback', 'beep');
pause(10)
close(h);
```

Output 12.4.1:



In the foregoing example, the action to be taken, sounding a beep, was triggered by a button press specified within the `uicontrol` definition. Typically, callback operations are more complex, so the actions to be taken are placed in local or nested functions rather than directly in the `uicontrol` callback definition. The following code illustrates this point. It yields the same results as the previous code, but the `uicontrol` definition provides a *pointer* to the callback routine, `@beepcallback`. The `@` symbol is a pointer to a function. If the command `beep` means “sound a beep,” the command `@beepcallback` means “do the operation specified by the function `beepcallback`,” which (if the function content matches the name) will also sound a beep.


Code 12.4.2:

```

function main
beepcount = 0;
h = figure;
hpb = uicontrol('Style','pushbutton',...
    'String','Click to Beep!',...
    'Position',[150 100 200 200],...
    'Callback',@beepcallback);
pause(10)
close(h);
msgbox(sprintf('Counted %d beeps!',beepcount));
return

    function beepcallback(source,eventdata)
        beep
        beepcount = beepcount + 1;
        return
    end
end

```

Output 12.4.2:

In the code just given, the callback function `beepcallback` is never called directly by the main function, which might lead you to ask, “How, then, does it get invoked?” Whenever the button is pressed, it’s “as if” the *button* calls the callback function `beepcallback` directly. The callback function has two input arguments, *source* and

`eventdata` that convey useful information about what happened. In this example, all the callback function needs to know is that the button was pressed, so the input arguments can be ignored. Any other operations needed for the program, such as counting the beeps, can be executed in the callback function. In this instance, the callback function is a nested function, so common variables such as `beepcount` are available both to the main and nested functions.

In the next example, we illustrate this concept via a “front-end” interface for a program that was shown earlier. The program relied on command-line interaction with the user to determine the number of days in any month given the month and year specified by the user (Code 7.3.6).

Here’s how the new program (Code 12.4.3) works. The main function first initializes the variables `month` and `year` to default values so each will have a valid value in case the user does not change it and so the variables will be visible to all nested functions. Then the figure is created. Two popup menus are installed at convenient locations in the window, one for the month and one for the year. Showing two popup menus provides us with a way of showing you two ways to define the contents of a popup menu. One is by passing a cell array of strings (as we do for the months). The other is by providing a single string with the items delimited by the vertical bar (|) character (as we do for the year). The definition of each interface element includes a pointer to its own callback function (e.g., `@monthcallback`) later in the program. A text field is installed above the popup menus for the user’s guidance. Because there is no action associated with a static text field, it needs no callback routine. The pushbutton is installed with a pointer to its callback routine (`gobuttoncallback`). Next, the main function waits for the figure window to close via the command `uiwait (handletothefigure)`.

All other program operations are initiated by the callback routines when the user clicks on the various interface elements. Whenever the user chooses a new month or year in the popup menu, the corresponding variable is set by the callback routine assigned to that popup window. When the `gobutton` is pressed, its callback routine, `gobuttoncallback`, first uses a `msgbox` to inform the user of the month and year to be computed, then waits, via `uiwait (hmsg)`, for the user to dismiss the message dialog, and then closes the figure by `close (handletothefigure)`. Closing the figure deletes its handle, so the condition for which the main program has been waiting, `uiwait (handletothefigure)`, has been met.

Code 12.4.3:

```
function main;
% Initialize variables common to main and nested subfunctions
month = 'January';
year = 2001;
% Open the figure for the interface
handletothefigure = figure;

% Install a popup menu for the months
monthstrings = {
```

```

        'Month'
        'January'
        'February'
        'March'
        'April'
        'May'
        'June'
        'July'
        'August'
        'September'
        'October'
        'November'
        'December'});
hmonth = uicontrol('Style', 'popupmenu',...
    'String',monthstrings,...
    'Position', [120 320 100 20],...
    'Callback', @monthcallback);

% Install a popup menu for the years
hyear = uicontrol('Style', 'popupmenu',...
    'String',...
    'Year'|2008|2009|2010|2011|2012|2013|2014|2015|2016|2017|2018',...
    'Position', [220 320 100 20],...
    'Callback', @yearcallback);

% Install an informative text field for the popup controls
uicontrol('Style','text','String','Pick a month and a year:',...
    'Position', [120,360,200,15]);

% Install a GO button
hgobutton = uicontrol('Style','pushbutton',...
    'String','Look up days in the month',...
    'Position', [120 120 200 40],...
    'Callback', @gobuttoncallback);

% Now just wait for the user to finish (when the window closes);
uiwait(handles.tothefigure)
return % from main

% Callback routines in nested functions:
function monthcallback(source,eventdata)
    mylist = (get(source,'String'));
    myitem = (get(source,'Value'));
    month = char(mylist(myitem));
end


```

```
function yearcallback(source,eventdata)
    mylist = (get(source,'String'));
    myitem = (get(source,'Value'));
    year = str2num(mylist(myitem,:))
end

function gobuttoncallback(source,eventdata)
    hmsg = msgbox(sprintf(...
        'Will compute days for %s, %d\n\n',month, year));
    uiwait(hmsg)
% When user presses the "go" button, the computation from
% Code 7.3.6 would be executed here, to return the results
% (code to be inserted).
    close(handletothefigure)
end

end %function main
```

Output 12.4.3:



12.5 Prototyping User Interfaces Using GUIDE

As noted above, GUIDE is MathWorks's drag-and-drop utility for GUI construction. GUIDE facilitates the placement of user controls in a window (saved as a .fig file), and automatically generates code for those user controls (saved as a .m file with the same name as the .fig file). An interface is built in GUIDE by dragging icons onto a representation of the eventual interface window, so that the placement, size, and contents of each of the elements can be adjusted. GUIDE then automatically generates the program file for that interface, which includes placeholders for the callback routines needed for the interface


elements the programmer has installed in the window. To complete the program design, the programmer fills in the details in the callback routines so each callback routine responds appropriately to the event that called it.

In our view, GUIDE is most useful for sophisticated programmers building complex interfaces. The needs of most readers of this text will be more easily met by the other two methods of interface construction outlined above. In fact, we encourage you to postpone your use of GUIDE until after you have done some explicit programming of interfaces using uicontrols (see Section 12.4). This will help you in your eventual exploration of GUIDE. When you feel ready to be “GUIDED,” you can watch GUIDE’s video tutorials, easily found via a search for “MathWorks GUIDE tutorial.” Search the MATLAB documentation for “A Working GUI With Many Components” for an example of using GUIDE to implement a number of different uicontrols in the same window.

These cautionary remarks having been made, we do want to provide you with one example of how GUIDE might be used to implement a program like Code 12.4.3. On the way to that demonstration, remember that a program with a GUI has two components: a figure (.fig) file that contains the interface elements, and a program (.m) file that contains the operational code that implements the interface functions. GUIDE automatically generates both the figure (as a .fig file) and the operational code (as an .m file with the same name as the .fig file), making it possible to rapidly prototype complex interfaces.

To work through the example, give the command `guide` in the Command window, and in the resulting dialog box, select Blank GUI (the default). After you have installed the indicated controls by selecting them in the left-margin menu, their placement should resemble the following.

Output 12.5.1:



Each of the controls in the windows is analogous to the corresponding control constructed in Code 12.4.3. Both windows have four interface elements: a text field, two pop-up menus, and a pushbutton. Using GUIDE, each of the elements was put into the figure by selecting its style from the left-hand palette and then inserting the element into the figure at the desired location and size, using familiar drag-and-drop techniques. GUIDE allows you to

move and resize each of the elements of the interface, and automatically generates the “first draft” of a `.m` file for the program, which already has an outline of the callback routines for each interface element. After placing the elements, you can label them, and provide other necessary information, using the Property Inspector, by right-clicking on each interface element or selecting Property Inspector from the View menu of `untitled.fig`. For example, in this case, the static text and second pop-up menu still have their default names, but we have already set the String variable of the first pop-up menu to the days of the month, and changed the String field of the button from its default (“Push Button”) to “Compute Date.”

When this figure is saved (as `Guideexample_12_5_1.fig`), in addition to the `.fig` file that captures the figure, a new `.m` file, `Guideexample_12_5_1.m`, is automatically generated with a main function and a nested callback function for each the interface elements. If all is well, the `.m` file will run without error, and the controls will seem to operate when you click on them. However, `.m` file won’t actually *do* anything because the operations of the callback functions have not yet been specified. It’s up to you to add to each callback function the code that will generate the operations that need to be performed. You can learn about these callback functions by examining the new `.m` file, where there will be one callback function for each operation you can perform in your new GUI. When you are done, the resulting callback functions will look similar to those of Code 12.4.3.

12.6 Recording User Interactions

Several of the GUI examples above have a common characteristic: More than one event may contribute to the program’s operation. In other words, how a program operates is determined not only by the program, but also by the behavior of the user. This “multi-responsiveness” is implemented through the use of callback routines. Another context where programs must respond to multiple, unpredictable events is in programs that behavioral scientists write to gather data such as reaction times.

Consider a simple reaction-time experiment in which you measure the time between presentation of an arithmetic problem and its solution. Code 6.4.1 used `tíc` and `tóç` for this purpose. Here is that program again, with its code number updated to set the stage for its amendment.

Code 12.6.1:

```
commandwindow
tic
response = input('What is five plus the square root of 64?')
Reaction_Time = toc
```

Output 12.6.1:

```
response =
    13
Reaction_Time =
    2.2859
```

The foregoing program provides an estimate of the time between presentation of the question and depression of the Enter key after all the keys used to type in the answer have been

pressed. But what if you need a precise measure of a single key-press response latency and that key press happens not to be the Enter key? Here is one approach.

The following example begins with a pause (inter-trial-interval) of random duration. The unpredictability of the delay helps discourage anticipatory responses. After the pause is over, an imperative “go” stimulus is presented. The `waitforbuttonpress` function reports either a button press or a key press, returning a value of 0 or 1, respectively, and reports the reaction time. Which key was pressed can be returned by `get(gcf, 'CurrentCharacter')`.

Code 12.6.2:

```
function RunATrial
figure(1);clf
pause(2+randi(4)/2)
text(.5,.5,'go','fontsize',32)
axis off
tic
waitforbuttonpress
reactiontime = toc
keypressed = get(gcf, 'CurrentCharacter')
close(1)
end
```

Output 12.6.2:

```
reactiontime =
    0.3626
keypressed =
x
```

The above code would respond to any key (or mouse click), and wait indefinitely for it to be pressed. If you were interested in detecting a particular key, you could examine the `keypressed` variable. You might also wish to put a time limit on the response, both to encourage fast responding and to move on in the event a participant dozes off.

The following program reports a reaction time for a press of the “g” key, and reports an error for any other key. To detect the key, it uses a callback function available in every MATLAB window. The callback function sets the figure’s `keypressfcn` (one of the figure’s attributes that can be read by `get` and modified by `set`) to point to a callback function that we call `gotAKey`, using the function pointer `@` operator.

This callback function (functionally similar to those used in Code 12.4.3) is executed when four conditions are met: the figure’s `keypressfcn` has been set to `@gotAKey`; the main function is inactive (not actually computing); the window is active (i.e., `frontmost`); and (of course) that a key is pressed. So, the main program activates the window, sets the `keypressfcn` to `@gotAKey`, and pauses for 3 seconds to allow a response, during which any key press will activate the callback function.

When the callback function executes, the key most recently hit can be retrieved from a field of the second variable passed to the callback routine by the key press, `event.Character`.

The program closes the figure either when a key press occurs or the pause has elapsed (the window could stay open, if there were more trials to come). At the end of 3 seconds (whether or not there has been a key press) the main function resumes. The callback function is written as a nested function to facilitate its sharing variables with the main function, `RunOneTrial`. The output shows three runs of the program from the Command window.

Code 12.6.3:

```
function RunOneTrial
myfigure = figure(1);clf
pause(2+randi(4)/2)
text(.5,.5,'Press!','fontsize',32)
axis off
reactiontime = [];
tic
set(myfigure,'KeyPressFcn',@gotAKey)
timedout = true;
pause(3)
close all;
% Other computation, such as recording the data
if timedout
    disp('timed out')
elseif correct
    fprintf('Reaction Time = %f\n',reactiontime)
else
    disp('It was an error')
    beep
end

    function gotAKey(src,event)
        timedout = false;
        if strcmp(event.Character,'g');
            reactiontime = toc;
            correct = true;
            set(myfigure,'KeyPressFcn',[])
        else
            correct = false;
        end
    end
end
```

Output 12.6.3:

```
>> Code_12_6_3
Reaction Time = 0.749595
>> Code_12_6_3
```

```

| timed out
| >> Code_12_6_3
| It was an error

```

The foregoing program detects correct keys, but it has a drawback. The program will not actually continue until the `pause` has completed, even if a response happens right after the onset of the stimulus. Is there some way to get the program to go on immediately if there is a response before 3 seconds have passed? You can use the `timer` function, which was introduced in Section 11.10. The `timer` function has a callback function, similar to the `uicontrols` described above. The timer is initiated by defining a handle to it (`mytimer`) and its callback function (`timercallback`), defining the action to be executed at the end of the interval using `@timercallback`. Then `start(mytimer)` starts the timer, and the timer runs until either the time is up, as reported by the execution of the callback function, or the timer is stopped by `stop(mytimer)` when a response occurs before time is up.

The definition of the timer includes a `startdelay` variable, which in this case is set to execute the callback function after a delay of 3 seconds. The `wait(mytimer)` command functions like the `pause` command, but can be interrupted by `stop(mytimer)` if a key is pressed, unlike `pause`. As soon as a response is detected, the program reports the reaction time of the trial or gives immediate error feedback (a beep for the wrong key or timeout) if needed. (No output is shown below because it would be similar to Output 12.6.3.)

Code 12.6.4:

```

function RunOneTrial
myfigure = figure(1);clf
pause(2+randi(4)/2)
text(.5,.5,'Press!','fontsize',32)
axis off
reactiontime = [];
tic
set(myfigure,'KeyPressFcn',@gotAKey)
timedout = false;
mytimer = timer('TimerFcn', @timercallback, 'startDelay', 3);
start(mytimer)
wait(mytimer)
if timedout
    disp('timed out')
    beep
elseif correctresponse
    fprintf('Reaction time = %f\n',reactiontime);
else
    disp('error')
    beep
end
close(myfigure)
% ... Other computation, such as recording the data
return

```



```

function gotAKey(src,event)
    correctresponse = false;
    if strcmp(event.Character,'g')
        reactiontime = toc;
        correctresponse = true;
    end
    timedout = false;
    stop(mytimer)
    set(myfigure,'KeyPressFcn',[])
end

function timercallback(src,event)
    timedout = true;
end

end

```

Let's put this information to use in a slightly more complicated experiment. We are interested in replicating the so-called Simon effect. Here, reaction times tend to be shorter when a stimulus and response are spatially compatible than when a stimulus and response are spatially *incompatible*, even if the spatial incompatibility is strictly irrelevant to the stimulus identification. For a review, see Lu and Proctor (1995).

We start with a variant of a discrimination experiment in which one of two symbols, L or R, appears on the screen on each trial. The L calls for a left response (the “a” key), whereas the R calls for a right response (the “;” key). If the Simon effect were replicated, responses to the letter L, which calls for a left-hand response, would be faster when L is shown on the left side of the screen than when L is shown on the right side of the screen, and vice versa for responses to the letter R.

In the experiment that follows, we have four trial types: L on the left, R on the right; L on the right, and R on the left. The first two types use compatible stimulus-response mappings. The second two use incompatible mappings. The program has several sections.

The `FilesSetup` section opens two files for output, one for summary data and one for text-based trial-by-trial data. Then it puts a header line in the data file. The `SetScreen` section makes a window across the bottom of the computer monitor. This window contains a central fixation point. The imperative stimulus appears 10% or 90% of the way across the window. The `DefineTrialTypes` section does two things. Using the 1×4 structure array `ttype`, it first specifies the conditions for each trial type. `ttype(1)`, for example, represents trials in which the L appears on the left (L), so these are Compatible (C) trials. `ttype(2)` represents trials in which the L appears on the right (R), so these are Incompatible (I) trials. Each trial type also has a field, `ttype(1:4).RT`, reserved for the reaction times to be acquired in that condition, and a counter, `ttype(1:4).error`, to count errors. The `InitializeData` section assigns an empty array to `ttype(1:4).RT`, and zero to `ttype(1:4).error`. All of these fields are initialized using the `deal` command (see Section 7.4). Finally, the `RuntheTrials` section

presents 32 trials, eight of each type, in random order, using the variable `typenum` to control which type of trial appears on each. If the response is correct, the reaction time is appended to the `ttype(typenum).RT` array, whereas if the response is incorrect, 1 is added to `ttype(typenum).error`. In addition, all relevant data from each trial are written to the `raw-data.txt` file.

Code 12.6.5:

```
function SimonDemo;
clc
clear
close all;
sinit = input('Subject's initials: ','s');
outfilename = ['SimonData_' sinit];
rawdataoutfilename = strrep(outfilename, '_', '_Rawdata_');
rawdataoutfilename = strcat(rawdataoutfilename, '.txt');
rawdatafile = fopen(rawdataoutfilename, 'w');
fprintf(rawdatafile, 'Trial\tside\tstim\tcomp\tKey\tResp.\tRT\n');
screensize = get(0, 'screensize');
% SetScreen
hfig = figure('position', [0 0 screensize(3) 200], 'color', [1 1 1]);

% DefineTrialTypes
[ttype(1:4).side] = deal('L', 'R', 'L', 'R');
[ttype(1:4).stim] = deal('L', 'L', 'R', 'R');
[ttype(1:4).comp] = deal('C', 'I', 'I', 'C');

% InitializeData.
[ttype(1:4).RT] = deal([]);
[ttype(1:4).error] = deal(0);
%Run 8 blocks of the four types in random order (32 in all);
trialnumber = 0;
for blocknumber = 1:8
    for typenum = randperm(4);
        trialnumber = trialnumber + 1;
        pause(2)
        hfix = text(.5, .5, '+', 'fontSize', stimulusfontsize);
        axis off
        set(gca, 'position', [0 0 1 1])
        pause(1)
        % Run the trial
        if ttype(typenum).side == 'L'
            stimposition = .1;
        else
            stimposition = .9;
        end
        hstim = text(.1, .5, ttype(typenum).stim, 'fontSize', ...
            72, 'fontweight', 'bold');
```

```
        tic
        waitforbuttonpress
        % Record the response
        thisRT = toc;
        thechar = get(gcf, 'CurrentCharacter');
        delete([hfix hstim]);
        switch thechar
            case 'a'
                thisResp = 'L';
            case ';'
                thisResp = 'R';
            otherwise
                thisResp = 'X'; % illegal key
        end
        if ttype(typenum).stim == thisResp
            ttype(typenum).RT = [ttype(typenum).RT thisRT];

fprintf(rawdatafile, '%2d\t%s\t%s\t%s\t%s\tcorrect\t%5.2f\n', ...
        trialnumber, ...
        ttype(typenum).side, ...
        ttype(typenum).stim, ...
        ttype(typenum).comp, thisResp, thisRT);
    else
        ttype(typenum).error = ttype(typenum).error + 1;
        beep

fprintf(rawdatafile, '%2d\t%s\t%s\t%s\t%s\terror\t%5.2f\n', ...
        trialnumber, ...
        ttype(typenum).side, ...
        ttype(typenum).stim, ...
        ttype(typenum).comp, thisResp, thisRT);
    end
end
end
save(outfilename, 'ttype');
fclose(rawdatafile);
```

If you declare your initials to be E. F., your participation in this experiment will generate two output files, `SimonData_ef.mat` and `SimonData_Rawdata_ef.txt`. The `.mat` file can be examined directly (see Code 12.6.6), or you could write a short program to report summary data (mean reaction time for each of the conditions of the experiment).

Code 12.6.6:

```
load SimonData_ef.mat
ttype
ttype1data = ttype(1)
```

Output 12.6.6:

```

ttype =
1x4 struct array with fields:
    side
    stim
    comp
    RT
    error
typedata =
    side: 'L'
    stim: 'L'
    comp: 'C'
    RT: [0.4938 0.6327 0.6572 0.8561 0.6182 0.7189 0.8305]
    error: 1

```

The other output file, `SimonData_Rawdata_ef.txt`, contains the results of all the trials in order, which might be needed, for example, if you were analyzing trial-by-trial dependencies in performance. Again, you could easily write a program to compute the mean reaction time for compatible and incompatible trials.

Code 12.6.7:

```
| type SimonData_Rawdata_ef.txt
```

Output 12.6.7:

Trial	side	stim	comp	Key	Resp.	RT
1	L	L	C	R	error	0.73
2	R	R	C	R	correct	0.79
3	L	R	I	R	correct	0.54
4	R	L	I	L	correct	0.51
5	L	R	I	R	correct	0.44
6	L	L	C	L	correct	0.49
7	R	L	I	R	error	0.39
8	R	R	C	R	correct	0.68
...data from 24 more trials not shown						

12.7 Practicing Interfaces and Interactions

Try your hand at the following exercises, using only the methods introduced so far in this book or in information given in the problems themselves.

Problem 12.7.1:

Finish the interface implementing the callback routines in `Guideexample_12_5_1.m`, and `Guideexample_12_5_1.fig`, which you will generate using the example of `Output_12_5_1` and `GUIDE`, by finding the callback routines in the code generated by `GUIDE` and inserting the appropriate code to respond to the ‘users’ actions.

Problem 12.7.2:

Combine Code 7.3.9 with Code 12.4.3, to make a GUI-based program for computing the number of days in any month. Provide the output (e.g., 'February 2008 has 28 days ') using an appropriate user control, in the same window as the input.

Problem 12.7.3:

Using `uicontrol` or `GUIDE`, devise an appropriate user interface for a program that you have previously written.

Problem 12.7.4:

Write a program to analyze the file `SimonData_Rawdata_ef.mat` from the website (or your own data file from running Code 12.6.5, or a similar tab-delimited text file of your choosing) and report the reaction times for the four conditions. Make the output suitable for transfer to a spreadsheet or statistics program.

Problem 12.7.5:

Write a program to analyze the file `SimonData_Rawdata_ef.txt` from the website (or your own data file from running Code 12.6.5, or a similar tab-delimited text file of your choosing) and report the reaction times for the four conditions. Make the output suitable for transfer to a spreadsheet or statistics program. Do the results agree exactly with those of Problem 12.7.4? Should they? (See Section 7.6 to help you get started.)

Problem 12.7.6:

Present a brief tone of moderate intensity, and allow the user to raise or lower the intensity of the tone using the up-arrow and down-arrow keys of the keyboard. Instruct the user to lower the intensity for the next trial if she hears it, and raise the intensity if she does not. Plot the psychophysical absolute threshold determined in this way, using a staircase graph. Make it possible to vary the frequency of the test stimulus. If you’d like, use a patch of gray on a dark gray background, instead, and plot the difference threshold as the patch is adjusted brighter and darker. Vary the brightness of the background on different trials, and determine the ratio between the difference threshold and the brightness of the background. Consider what precautions should be taken to ensure that performance is not affected by non-sensory factors, such as stimulus sequence.

13. Psychtoolbox

This chapter covers the following topics:

- 13.1 Introducing Psychtoolbox
- 13.2 Installing Psychtoolbox
- 13.3 Writing a simple Psychtoolbox program
- 13.4 Using Psychtoolbox documentation
- 13.5 Changing fonts and font sizes
- 13.6 Adding shapes to a display
- 13.7 Adding textures and images to a display
- 13.8 Displaying stimuli sequentially with precise timing
- 13.9 Collecting keyboard input
- 13.10 Monitoring keyboard input while doing other things
- 13.11 Collecting a response string
- 13.12 Collecting mouse data
- 13.13 Creating an animation with moving dots
- 13.14 Making things transparent
- 13.15 Testing the Simon effect with Psychtoolbox
- 13.16 Exploring Psychtoolbox further
- 13.17 Recovering from Psychtoolbox program crashes and infinite loops
- 13.18 Problems

The commands that are introduced and the sections in which they are premiered are:

GetSecs	(13.2)
ScreenTest	(13.2)
SetupPsychtoolbox	(13.2)
sca	(13.3)
Screen	(13.3)
Screen('DrawText')	(13.3)
Screen('Flip')	(13.3)
Screen('Openwindow')	(13.3)
WaitSecs	(13.3)
Screen('Preference')	(13.5)
Screen('TextFont')	(13.5)
Screen('TextSize')	(13.5)
Screen('FrameOval')	(13.6)
Screen('DrawTexture')	(13.7)
Screen('MakeTexture')	(13.7)

<code>Screen('FillOval')</code>	(13.8)
<code>Screen('GetFlipInterval')</code>	(13.8)
<code>KbCheck</code>	(13.9)
<code>KbName</code>	(13.9)
<code>ListenChar</code>	(13.10)
<code>GetMouse</code>	(13.12)
<code>HideCursor</code>	(13.12)
<code>set(gca, 'YDir', 'reverse')</code>	(13.12)
<code>SetMouse</code>	(13.12)
<code>ShowCursor</code>	(13.12)
<code>Screen('DrawDots')</code>	(13.13)
<code>KbPressWait</code>	(13.14)
<code>Screen('BlendFunction')</code>	(13.14)
<code>Beeper</code>	(13.15)
<code>TextBounds</code>	(13.15)

13.1 Introducing Psychtoolbox

As powerful as MATLAB is, it has limitations. Some of the limitations make it difficult to use MATLAB to full advantage in behavioral science, especially when running behavioral experiments that require high spatial or temporal resolution. If you are a behavioral scientist, established or aspiring, you should be aware of these limits. They have been discussed by Plant and Turner (2009) and Plant and Quinlan (2013).

Happily, it is possible to use MATLAB *toolboxes* to get around the limitations. Toolboxes, in general, are suites of programs that are designed to serve special purposes. Some toolboxes have been developed by the MathWorks (the company behind MATLAB), but those toolboxes have not been developed specifically for behavioral science needs. Other toolboxes have been developed by others, not directly associated with the MathWorks for behavioral science. These toolboxes are free. One is MatTap, short for MATLAB Timing Analysis Package (Elliott, Welchman, & Wing, 2009; see www.snipurl.com/mattap to download the toolbox). Other toolboxes for presenting stimuli and recording responses can be found at www.hans.strasburger.de/psy_soft.html.

One of the freely available toolboxes for behavioral science research is Psychtoolbox (short for Psychophysics Toolbox). Psychtoolbox is now used so widely by behavioral scientists that we devote a full chapter to it in this book. We believe you will be able to use Psychtoolbox more easily having gone through this chapter than striking out on your own. No other source that we are aware of gives as much information

about the nuts and bolts of Psychtoolbox as does this chapter. If you present research that used Psychtoolbox, please follow the accepted practice of citing Brainard (1997), Pelli (1997), and Kleiner et al. (2007).

Why is Psychtoolbox so popular? There are four reasons:

1. **Greater video speed.** Psychtoolbox communicates directly with the computer's video hardware and so permits extremely rapid drawing of complex stimuli.
2. **Greater temporal reliability.** Presenting stimuli and collecting data with your computer often require very high temporal precision. Because Psychtoolbox communicates directly with the computer's hardware, it enables highly reliable timing of stimulus onsets and user inputs. This reliability typically allows for more accurate timing than is possible with MATLAB alone. Psychtoolbox compares favorably with, and may even surpass, commercial software packages in this regard.
3. **Hardware flexibility.** Psychtoolbox enables communication with other devices that provide more ways to gather data.
4. **Compatibility.** Like MATLAB, Psychtoolbox works on Microsoft Windows, Linux, and Mac platforms. Psychtoolbox is also compatible with a free program called OCTAVE (www.gnu.org/software/octave/), which performs many but not all of the functions of MATLAB.

13.2 Installing Psychtoolbox

To use Psychtoolbox, you will need to install it or have it installed on your computer. (In this chapter, we assume Psychtoolbox Version 3.) To find out how to install Psychtoolbox, consult the Psychtoolbox website (<http://psychtoolbox.org>) and navigate to the installation webpage. We will not repeat the instructions here because they vary depending on the operating system. However, there is a shortcut to installation if you have access to another machine that has Psychtoolbox installed on it, it has the same hardware, has the same operating system, and has the same version of MATLAB. In that case, you can copy the complete `Psychtoolbox` folder from one machine to the other. Then, on the new machine, you can change your MATLAB directory to the new location of the Psychtoolbox folder and run the `SetupPsychtoolbox` function in that folder. You should have administrator privileges for this task because the installation script will need to modify some settings in your MATLAB folder.

For Psychtoolbox to work well on your computer, there should be a dedicated graphics card that allows it to display visual stimuli rapidly. Most modern laptops of either the Windows or Macintosh variety have such a card. If you are running Microsoft Windows, you will need to check that you have installed the dedicated graphics drivers for your video card.

It's important to note that Psychtoolbox runs with better timing on systems that have only a single monitor (or two mirrored, functionally identical, monitors) than on systems with a desktop that spans multiple monitors. To find out more about different hardware configurations, consult a help file called "Synctrouble" on the Psychtoolbox website.

Once Psychtoolbox has been installed, you can test whether MATLAB can access it via the following command:

Code 13.2.1:

```
| format long;  
| Timenow = GetSecs
```

Output 13.2.1:

```
| Timenow =  
|          8.007085074686600e+04
```

If MATLAB returns an error message, something has obviously gone wrong. One possibility is that the `SetupPsychtoolbox` command didn't have permission to change MATLAB's path settings. Be sure to run the setup function from an account with administrator privileges and be prepared to enter the account's password.

Speaking of things that can go wrong, once you have successfully installed Psychtoolbox, there is some chance your program might crash, leaving your computer under the control of Psychtoolbox but without any clear way to exit. If this happens, don't panic. You can restore control of your computer via a few keystrokes that will be described in Section 13.17.

Returning to Code 13.2.1, we assigned the returned value from `GetSecs` to a variable whose name we chose, `Timenow`. The function `GetSecs` is a reserved term in Psychtoolbox. `GetSecs` returns a time value in seconds. You can use this value to measure such things as how long it took someone to respond to a stimulus.

Once you have Psychtoolbox running, you can determine how well it operates on your system using a function called `ScreenTest`.

Code 13.2.2:

```
| ScreenTest
```

When you run `ScreenTest`, you will see the screen go blank, and then you will see the phrase, "Welcome to Psychtoolbox." After this, you should see the MATLAB Command window, where you should see something like the following (depending on your operating system):

Output 13.2.2:

```
| ***** ScreenTest: Testing Screen 0 *****  
  
| PTB-INFO: This is Psychtoolbox-3 for Apple OS X, under  
| Matlab 64-Bit (Version 3.0.11 - Build date: Jul 23 2013).
```

```
PTB-INFO: Type 'PsychtoolboxVersion' for more detailed
version information.
...
```

The output is too long to display here, but it is worth reading on your computer, for it may give information about incompatibilities with your graphics hardware and Psychtoolbox. You may see a warning about “SYNCHRONIZATION TROUBLE” in the MATLAB command window, or a big red flashing warning sign after the screen goes blank. This problem is probably related to your computer’s graphics hardware. Your computer has not been damaged if you get either of these messages. These warnings mean that the timing may be a bit imprecise on the computer you are using. In some cases, it is necessary to restart MATLAB to resolve synchronization errors.

Fortunately, there are online resources to help with problems like these. At the Psychtoolbox website (<http://psychtoolbox.org>), you can find links to important sources, including answers to Frequently Asked Questions (FAQs) and a support forum that is full of answers to problems of this sort. On the off chance that you can’t find a solution, you can post a question in the forum and someone in the community—a remarkably generous group—will probably respond quickly. You may choose to develop Psychtoolbox programs on computers that have synchronization issues, but the data collection should be performed on systems that run Psychtoolbox without such warnings. If you wish to set up a computer just for development, there is a way to disable synchronization tests. Psychtoolbox tells you how to do this in the warning message in the command window.

13.3 Writing a Simple Psychtoolbox Program

To help you begin programming with Psychtoolbox, we invite you to write a simple Psychtoolbox program that prints the phrase “Hello World!” Enter the following code in the MATLAB Editor and then run it.

Code 13.3.1:

```
mywindow = Screen('OpenWindow',0);
Screen('DrawText', mywindow, ...
    'Hello World!',200,100,[0,0,0]);
Screen('Flip',mywindow);
WaitSecs(1);
sca
```

If all goes well, the screen will switch to white and will show you a *Welcome* message, followed by this message in a font that may or may not match the one below.

Output 13.3.1:

```
| Hello World!
```

When the screen clears, you will be brought back to your desktop, where you should again see the diagnostic text you encountered in Output 13.2.2.

Code 13.3.1 deserves a few extra comments. First, Psychtoolbox uses a function called `Screen` for many of its operations. You can specify what `Screen` does by passing a string to it. The string passed to `Screen` in Code 13.3.1 is `'OpenWindow'`. This tells Psychtoolbox to take control of your computer screen. The next argument, the number 0, tells `Screen` which monitor to use, in this case your primary monitor. If you want to use a secondary monitor, you can use that monitor's number (e.g., 1) instead. Beware that Psychtoolbox can behave strangely on computers with multiple monitors, as mentioned in Section 13.2.

Second, when `Screen` initiates a Psychtoolbox session, it returns a value called a *window pointer*. You should assign this value to a variable of your choosing for later reference. In Code 13.3.1, the variable is called `mywindow`, which is used in the second and third lines of the program. There is nothing special about this name. We could have used `banana_peel`. The variable becomes a window pointer because of where it appears in the function result, not because of its verbal label.

Third, Code 13.3.1 uses `Screen` to draw text via `DrawText`. The `DrawText` command requires additional arguments to determine what will actually be drawn on the screen. The first argument, as already mentioned, is the window pointer, which specifies the window in which text will appear. The next argument defines the characters to display. In this case, the characters comprise the string `Hello World!` The next two arguments specify where the text will be shown. These two arguments specify the horizontal and vertical coordinates, in screen pixels, of the upper left corner of the first character. Psychtoolbox defines the upper left corner of the screen to be the coordinates (0, 0), so the pixel scale runs to the right and down from this point. Given this convention, the second line of Code 13.3.1 tells Psychtoolbox to place the upper left corner of the `H` in `Hello` 200 pixels from the left of the screen and 100 pixels down from the top.

The final argument for `DrawText` is a vector with three values specifying the levels of red, green, and blue, respectively. On most computers, in Psychtoolbox the color scale goes from 0 to 255 rather than from 0 to 1, which is the default in MATLAB (see Sections 9.5 and 10.2). The 256 possible values (0 through 255) for each color specify the color's intensity. If all of the colors get zero energy, then none of them is brightened and the color is black, as in the code above. Alternatively, if all of the colors get maximum energy, [255, 255, 255], all of them are brightened fully and the color is white. By using different numbers, you can dip your computer paintbrush into a rich palette. The richness of this palette is quantifiable as $256^3 = 16,777,216$ possible hues. Specialized hardware can provide an even larger color palette which is useful when using grayscale stimuli to study the visual system.

In our discussion of Code 13.3.1, we have gotten to the end of the second line. If we stopped the program here, nothing would be displayed on the screen. Why not? The reason is that Psychtoolbox allows you to complete all the drawing you want ahead of time, in a hidden window, before making that window visible, using a special command called `Flip`. The `Flip` command is used here in the fourth line of Code 13.3.1. Before the command is issued, the text that is drawn (if `DrawText` is used ahead of time) is prepared for showing, but it isn't actually shown until the "card" on which it is drawn is "flipped." Being able to draw before displaying lets you draw many stimuli before revealing all of them at once, without flicker, with a single command. Suffice it to say that this capability makes Psychtoolbox a good choice for vision experiments or other behavioral science projects in which several components of a complex visual stimulus need to be presented simultaneously. As you might expect, textual stimuli are not the only

ones you can show in this way. Non-textual visual stimuli (e.g. shapes and images) can also be shown, as discussed later in this chapter.

The fifth line of Code 13.3.1 causes MATLAB to wait for 1 second, using the `WaitSecs` command. This command is similar to MATLAB's `pause` command but is more precise. The timing accuracy of `WaitSecs` is less than 1 millisecond, whereas the timing accuracy of `pause` can be off by 10 milliseconds or more depending on the configuration of your computer.

The sixth and final line of Code 13.3.1 shuts down Psychtoolbox, closing the Psychtoolbox window and returning control to MATLAB, by use of the `sca` command. `sca` is shorthand for `Screen('CloseAll')`. You must shut down Psychtoolbox at the end of every Psychtoolbox program. If you don't, the program ends with Psychtoolbox still in control of your computer and you need to resort to the techniques in Section 13.17 to regain control. It's better to `sca` than to scamper through those hoops.

13.4 Using Psychtoolbox Documentation

There are many possible commands in Psychtoolbox and there are several ways to learn about them. Via the Internet, you can access an online list of Psychtoolbox functions (<http://docs.psychtoolbox.org>). In addition or instead, you can use MATLAB's familiar `help` command, as in the example below, which yields a helpful reply if Psychtoolbox is installed. The command elicits lengthy, but potentially informative text. We have omitted the output here but recommend that you seek such help.

Code 13.4.1:

```
| help Screen
```

The `help` command works for all of the basic Psychtoolbox commands, such as `Screen`, `GetSecs`, and `WaitSecs`.

If you want to get more information about the various operations that can be performed with the `Screen` function, you can use the `Screen` command to provide additional help by following it with the name of an operation, followed by a question mark, as shown below. The output from this command is lengthy. We just show some of it.

Code 13.4.2:

```
| Screen DrawText?
```

Output 13.4.2:

```
Usage:
[ newX, newY ] = Screen('DrawText', windowPtr, text [, x]
[, y] [, color] [, backgroundColor] [, yPositionIsBaseline]
[, swapTextDirection]);
```

```
Draw text. "text" may include Unicode characters (e.g.
Chinese). A standard MATLAB/Octave character text string
is interpreted according to Screen's current character
encoding setting
...
```

The full output lists all of the possible arguments that can be applied when you use the `DrawText` operation as well as a description of how those arguments can be used. Some of the arguments are enclosed in square brackets, `[]`. In this context, the square brackets do not denote a MATLAB array. Instead, they specify optional arguments. If you do not specify the optional arguments, a default value will be used, as mentioned in the documentation.

13.5 Changing Fonts and Font Sizes

The font and size of text can be controlled in Psychtoolbox, just as it can in MATLAB Figure windows. Here is a more elaborate version of Code 13.3.1 which includes commands to change the font and font size.

Code 13.5.1:

```
Screen('Preference', 'VisualDebugLevel', 1);
window = Screen('OpenWindow', 0);
Screen('TextSize', window, 50);
Screen('TextFont', window, 'Times');
Screen('DrawText', ...
    window, 'Hello World!', 100, 100, [0, 0, 0]);
Screen('Flip', window);
WaitSecs(1)
sca
```

Reading the program, it should be obvious which commands specify the font size (50) and font identity (Times). You can use `Screen` (refer back to 13.4.2) to learn more about how these commands work.

Code 13.5.1 introduces another command in the first line, whose purpose is less obvious. The string `'Preference'` changes the `VisualDebugLevel` settings in Psychtoolbox so the standard welcome screen is replaced with a black background.

You can also set the background color of the screen by providing an RGB (red, green, blue) color specification for the `OpenWindow` command, as in the code below.

Code 13.5.2:

```
window = Screen('OpenWindow', 0, [255, 0, 0]);
Screen('Flip', window);
WaitSecs(1)
sca
```

Before you run the program, can you tell what color will appear on the screen? Hint: If you can't, you might look embarrassed (red in the face).

13.6 Adding Shapes to a Display

Psychtoolbox provides commands for drawing shapes such as circles, squares, and other polygons. For example the command `FrameOval` draws an empty circle. `FillOval` draws a filled circle. The example below provides for an empty circle and text.

Code 13.6.1:

```
Screen('Preference', 'VisualDebugLevel', 1);
mywindow = Screen('OpenWindow', 0);
Screen('TextSize', mywindow, 50);
Screen('TextFont', mywindow, 'Times');
Screen('DrawText', mywindow, 'Hello World!', 100, 100, [0, 0, 0]);
Screen('FrameOval', mywindow, [200 0 200], [75 50 225 200], 5);
Screen('Flip', mywindow);
WaitSecs(1);
Screen('Flip', mywindow);
WaitSecs(1);
sca
```

Output 13.6.1:



Something that is not apparent from the static output shown here is that on your computer, the output appears for about one second and then disappears. Also not apparent in the output above, but apparent on the website (www.routledge.com/9780415535946), is that the circle on your screen will be purple. The purple circle was drawn with `FrameOval`. Note that your display may look slightly different because fonts vary from one computer platform to another.

Just as `DrawText` needs additional information about what and where to draw the text, `FrameOval` needs additional information about the oval to be drawn. As before, the second argument is the window pointer. The third argument specifies the color, with three values for red, green, and blue, chosen here to yield the color purple. The fourth argument specifies the oval's dimensions by indicating the left, top, right, and bottom edges of the oval, in screen pixels, starting from the upper left corner of the screen. The fifth argument is the thickness, in pixels, of the frame around the oval.

The `Flip` command is issued after the oval and text are prepared for presentation. Once the `Flip` command is issued, the compound image of the text and oval are shown and remain on the screen until the time specified in `WaitSecs` has transpired, whereupon Psychtoolbox performs the next commanded operation. In this case, that operation is clearing the screen and then waiting for another full second before shutting down Psychtoolbox with the `sca` command.

An important lesson from this example is that every time the `Flip` command is issued, the display is updated. If you haven't issued any new draw commands since the most recent `Flip`, the next time you `Flip`, the screen will go blank.

13.7 Adding Textures and Images to a Display

Psychtoolbox also lets you show images using elements known as *textures*. For Psychtoolbox, a texture is an object into which you can place an image that you would like to display rapidly. The image could be as small as a period or as large as a photograph. Once created, textures can be placed on the screen in any way you choose, and they can be displayed rapidly via `Flip`.

Here is an example of how to make a texture object and display it on the screen. You may recognize this image from Chapter 10.

Code 13.7.1:

```
imagedata = imread('lab_photo.jpg');  
window = Screen('OpenWindow', 0);  
TexturePointer = Screen('MakeTexture', window, imagedata);  
clear imagedata;  
Screen('DrawTexture', window, TexturePointer);  
Screen('Flip', window);  
WaitSecs(2);  
sca
```

Output 13.7.1:



To create a texture, you need a matrix of numbers, as you do with MATLAB's `image` function, and to define such a matrix, you can load an image using MATLAB's `imread`, as in Section 10.3. Next, this matrix is passed to the `MakeTexture` function, which converts the matrix into a texture and returns a pointer to that texture. As with the window pointer discussed in Section 13.3.1, you can assign that pointer to a variable with any name you choose. This pointer can then be used to place the texture on the screen via the `DrawTexture` command. Once you have created the texture, you no longer need the original matrix. You can clear the matrix by using the `clear` command, as shown in the code above. Clearing a no-longer-needed matrix is prudent, especially if your computer's memory is running low.

Why use textures? There are three reasons. First, if you are using Psychtoolbox to control the screen, MATLAB's `image` function will not work. Second, once you have created a texture, your computer can display that texture very rapidly. In fact, textures can be displayed so quickly that you can put many of them on the screen simultaneously, in only a few milliseconds. Third, once you have created a texture, it can be stretched, shrunk, or rotated, and placed anywhere on the screen with a single command. You can learn about these capabilities by typing `Screen DrawTexture?` at the command prompt.

The input to `MakeTexture` is typically an $X \times Y \times 3$ matrix in which the color of each pixel is specified as RGB (red, green, blue) energy levels. Shades of gray can be created by setting the levels of the three colors to be equal. Therefore, values of `[200 200 200]` for R, G, and B yield a light gray, whereas `[50 50 50]` yield a dark gray; see Section 10.4. If you are creating an image containing only shades of gray, the input to `MakeTexture` can be just an $X \times Y$ matrix, where the grayscale brightness of each pixel is specified by a value (0 to 255).

13.8 Displaying Stimuli Sequentially With Precise Timing

Psychtoolbox lets you display multiple stimuli sequentially. Code 13.8 illustrates how this can be done. The program allows for the sequential display of two circles. The first circle remains on the screen for about 1 second and then is replaced by the second circle. The output is not shown here. If you run the program, you will see a circle near the top of the screen followed by a circle beneath it.

Code 13.8.1:

```
Screen('Preference', 'VisualDebugLevel', 1);
window = Screen('OpenWindow', 0);
Screen('FillOval', window, [0, 200, 200], [200, 200, 250, 250]);
onsetTime1 = Screen('Flip', window);
WaitSecs(1);
Screen('FillOval', window, [0, 200, 200], [200, 300, 250, 350]);
onsetTime2 = Screen('Flip', window);
WaitSecs(2);
sca
```

In the text before Code 13.8.1, we used the word *about* when describing the duration of the first display. We used that term because we were satisfied with an approximate duration for

that stimulus. But what if you needed an exact duration, such as a duration of *exactly* 1 second? To determine precisely how long the stimulus is displayed, you can use another feature of `Flip`, the ability to return the precise time that the stimulus was sent to the monitor. In Code 13.8.1, we store these values in two variables, `onsetTime1` and `onsetTime2`, and we subtract the second from the first to find the difference:

Code 13.8.2:

```
| onsetTime2-onsetTime1
```

Output 13.8.2:

```
| ans =  
|      1.0188
```

Why isn't the answer exactly 1.000? The reason is that the display of the second stimulus started only after the 1 second wait time elapsed, and it took some additional time to draw the oval and then `Flip` the screen.

Here is another version of the same program but with a much tighter degree of control over the stimulus duration.

Code 13.8.3:

```
| Screen('Preference', 'VisualDebugLevel', 1);  
| window = Screen('OpenWindow', 0);  
| halfFlip = Screen('GetFlipInterval', window)/2;  
| Screen('FillOval', window, [0, 200, 200], [200, 200, 250, 250]);  
| onsetTime1 = Screen('Flip', window);  
| Screen('FillOval', window, [0, 200, 200], [200, 300, 250, 350]);  
| onsetTime2 = Screen('Flip', window, onsetTime1 + 1.0 - halfFlip);  
| WaitSecs(2);  
| myduration = onsetTime2 - onsetTime1  
| sca
```

Output 13.8.3:

```
| myduration =  
|      0.9997
```

Now if you compute the time lag between the two `onsetTime` variables, you will see that it is much closer to 1.0000 second (assuming Psychtoolbox can synchronize with your video card).

We used two features of Psychtoolbox to achieve this. The first is the `Screen` operation `GetFlipInterval`, which returns the time it takes for your computer monitor to flip from one display to another. This value is typically equal to 1 divided by your monitor's refresh rate, so if the refresh rate is 100 Hz, this value is .01 seconds. In this example, half the value returned by `GetFlipInterval` is stored in the variable `halfFlip`.

The second new feature is specification of the exact time at which the second `Flip` should begin. The specification is achieved by providing an optional third argument to `Flip`. The specified time is `onsetTime1 + 1.0 - halfFlip`, or 1 second more than the onset of the first dot, minus one half the flip duration. The subtraction at the end is necessary because the `Flip` command takes time to execute, so it needs to begin slightly before the critical moment.

13.9 Collecting Keyboard Input

Psychtoolbox provides a way of collecting information from the keyboard without pausing the program. To take advantage of this capability, use `KbCheck`. Type the following command and press the Return key.

Code 13.9.1:

```
| x = KbCheck(-1)
```

Output 13.9.1:

```
| x =  
1
```

The output of 1 (in other words `true`) indicates that at the moment `KbCheck` was called, some key was pressed. If the returned value was 0, then you released the Return key quickly enough for `KbCheck` to miss it. Type the command again and you should see a different result.

An important point to consider when running a program in Psychtoolbox or, for that matter, when running *any* program in MATLAB, is that any function called by the program – be it `KbCheck` or some other function – takes longer to execute the first time it is called than later. Whenever a function is first called by a program in MATLAB, MATLAB has to load the function into memory; see `help GetSecs` for details. This extra delay needs to be taken into account in the design of experiments requiring maximally precise timing. You can minimize the effect of the delay by including practice trials that use the same functions as the rest of your trials. If you do not want your subject’s practice trials to be objectively different from subsequent “real” trials, you can run your code through one mock trial before the subject arrives.

Returning to Code 13.9.1, the argument `(-1)` to `KbCheck` causes it to check all of your attached keyboards, for example in the case of a laptop with an attached USB keyboard.

Which key was pressed? You can find out by recording all three of the variables returned by `KbCheck`.

Code 13.9.2:

```
| [KeyIsDown secs keyCode] = KbCheck(-1);
```

This line of code lets you retrieve three values from `KbCheck` instead of one. As usual, you can assign these three values to variables with any name you choose. The first variable, here named `KeyIsDown`, is a Boolean (0 or 1) that represents whether any key was down,

as illustrated in 13.9.1. The second variable returns the time that `KbCheck` was executed, providing a time stamp similar to that returned by `GetSecs`. This value is useful for computing precise estimates of reaction time. The third variable is a vector of 0's and 1's indicating which key or keys were down. Psychtoolbox also gives you a way to translate that vector of 0's and 1's into key names with a function called `KbName`.

Code 13.9.3:

```
| KbName(keyCode)
```

Output 13.9.3:

```
| ans =  
| Return
```

Psychtoolbox confirms that the key you hit was `Return`.

Key names differ for Mac, Windows, and Linux, so before using `KbName`, it is advisable to issue this command:

Code 13.9.4:

```
| KbName('UnifyKeyNames')
```

This command changes the names of keys across different computer platforms so they all match. That way, a program you write on Windows, for example, will return the same key names when it is run on a Mac.

13.10 Monitoring Keyboard Input While Doing Other Things

Recall that when using “naked” MATLAB, pausing for keyboard input with `waitfor` `buttonpress` or `input` causes all else to stop. We described one way to overcome this limitation using Figure windows in Section 12.6. Psychtoolbox also has a solution for this problem that relies on `KbCheck` and `GetSecs`, as shown in the code below. Here, keyboard input is awaited for up to 5 seconds, after which the program ends. We use `KbCheck(-1)` to include all attached keyboards in the check.

Code 13.10.1:

```
| while(KbCheck(-1))  
| end  
  
| waitTime = 5;  
| nowTime = GetSecs;  
| endTime = nowTime + waitTime;  
| keyDown = 0;  
  
| ListenChar(2);  
| while(keyDown ==0) & (nowTime < endTime)  
|     keyDown = KbCheck(-1);
```

```

        nowTime = GetSecs;
    end

    if(keyDown ==1)
        'A key was pressed'
    else
        'Ran out of time'
    end
    ListenChar(0);
    sca

```

Code 13.10.1 begins with a short `while` loop that serves as a safeguard to ensure that no keys are pressed before it begins. The short `while` loop ends once `KbCheck` returns false (0). Then, after setting a few parameters, the program uses a second `while` loop to keep checking the status of the keyboard with `KbCheck`, but also to check how much time has passed using `GetSecs`. If either a key is pressed or 5 seconds have elapsed the `while` loop ends.

Code 13.10.1 introduces another command called `ListenChar`. `ListenChar` controls how the program responds to keyboard input. This command is useful when `KbCheck` is used to collect keystrokes and you don't want the keystrokes to show up in the MATLAB command window. Use `ListenChar(2)` to block keystrokes from going to the main MATLAB window (although `KbCheck` will still be able to detect them) and `ListenChar(0)` to remove the block. To see why this is useful, comment out the `ListenChar(2)` line and run the program again. You will now be able to see the key you pressed appear in the MATLAB command window.

An important caveat about `ListenChar(2)` is that if your program quits or crashes without running `ListenChar(0)`, keyboard input will still be blocked. Don't panic. You can resort to `ctrl-c` to cancel the key press blockade and return MATLAB to normal.

13.11 Collecting a Response String

You may want to collect more than one keystroke within a Psychtoolbox program. You can do this with `KbCheck` by using a loop. Here is an example of code that checks the value of `keyCode` until five keystrokes have been detected.

Code 13.11.1:

```

KbName('UnifyKeyNames');
Chars = 0;
maxChars = 5;
Response = [];

[KeyIsDown secs keyCode] = KbCheck(-1);

ListenChar(2)
while Chars < maxChars
    lastkeyCode = keyCode;

```

```
[KeyIsDown secs keyCode] = KbCheck(-1);

difference = keyCode - lastkeyCode;
keys = find(difference == 1);

for(i = 1: length(keys))
    if(Chars < 5)
        Chars = Chars + 1;
        Response = [Response KbName(keys(i))];
    end
end
end
ListenChar(0)
fprintf('The typed response was \n%s\n',Response)
```

Run this program and type any five lowercase letters you like. You should see them echoed back after the fifth keystroke. Remember, if the program crashes, `ListenChar` will still suppress the keyboard input and you will need to press `ctrl-c` to use the keyboard again.

It will be helpful for your understanding of Code 13.11.1 to step through it. The program uses a `while` loop to record each keystroke. To detect a new keystroke, the output of `KbCheck` on any one execution of the `while` loop is compared to the output from the previous execution. If this comparison reveals that an element of `keyCode` switches from 0 to 1, the program registers a new key press. This comparison is performed by subtracting the previous value of `keyCode` from the new value of `keyCode` and then storing the result in the variable named `difference`. In this resultant vector, any values of '1' indicate a new key press. To see why this works, imagine subtracting the row vector `[0 0 0 1]` from `[1 0 0 1]`. The result is the vector `[1 0 0 0]`. The presence of a 1 indicates that the two vectors differ. After the subtraction, the program uses MATLAB's `find` function to determine which key numbers were pressed. These key numbers are then passed into `KbName` to extract the names of the new keystrokes.

Now run the program again and type following string: 'ab cd'. You will see the following output:

Output 13.11.1:

```
| The typed response was
| abspacedcd
```

This will seem strange at first, but remember that `KbName` returns the name of each key, including the spacebar, whose name is `space`. Fortunately, this problem can be overcome. If you wish to use `KbCheck` to collect response data that contains spaces, simply write code that converts the string `space` to the character ' '. If you press keys that have more than one character on them (like the comma key), `KbName` will return both of the characters. For example, run the program again and type the string 'op[]\'

Output 13.11.2:

```
| The typed response was
| op[{}]\|
```

Note that there are several other functions associated with `KbCheck` that provide shortcuts for some of the functionality described here. You will see them listed at the end of the help documentation for `KbCheck`. Such helpful links are present at the end of most of the help files in Psychtoolbox.

13.12 Collecting Mouse Data

Psychtoolbox is useful for collecting mouse data. You can get such data easily with the `GetMouse` function. Here is a program that asks you to trace a circle with the mouse and measures the mouse position. The program turns off the typical mouse cursor and replaces it with a mouse cursor drawn by Psychtoolbox.

Code 13.12.1:

```
% Part One: Initialization
Screen('Preference', 'VisualDebugLevel', 1);
window = Screen('OpenWindow', 0);
Screen('TextSize', window, 24);
Screen('TextFont', window, 'Times');
circleradius = 150;
circlecenter = 400;
textX = 200;
textY = 200 ;
Cursorsize= 6; %how big our mouse cursor will be
mousedata = zeros(10000,2); %used to store mouse data points
sample = 0;
%move the mouse to a specific spot
SetMouse(circlecenter-circleradius, circlecenter, window);
HideCursor; %hide the existing mouse cursor

% Part Two: Mousewait
buttons = 1;
while any(buttons)
    [Mousex, Mousey, buttons] = GetMouse(window);
end

% Part Three: Collect Data
DesiredSampleRate = 10 %Number of samples per second
clear sampletime;
begintime = GetSecs;
nextsampletime = begintime;
while buttons(1) ~=0
    sample = sample + 1;
    xlocation = 0;
    lowerbound = circlecenter-circleradius;
    upperbound = circlecenter+circleradius;
```

```


Screen('DrawText',window', ['Trace the Circle '...
      'clockwise, then click the left mouse button'],...
      textX,textY,[0, 0, 0 ]);
Screen('FrameOval', window , [0, 0, 0 ],[lowerbound ...
      lowerbound upperbound upperbound],4);
Screen('FrameOval', window , [0, 0, 0] , ...
      [Mousex-Cursorsize, ...
      Mousey-Cursorsize,Mousex+Cursorsize, ...
      Mousey+Cursorsize],3);
Screen('Flip',window);
[Mousex,Mousey,buttons] = GetMouse(window);
mousedata(sample,1) = Mousex;
mousedata(sample,2) = Mousey;
sampletime(sample) = GetSecs;
nextsampletime = nextsampletime + 1/DesiredSampleRate;
while GetSecs < nextsampletime
end
end

% Part Four: Cleanup
endtime = GetSecs;
ElapsedTime = endtime - begintime
NumberOfSamples = sample
ActualSampleRate = 1/(ElapsedTime / NumberOfSamples)
mousedata = mousedata(1:sample,1:2);
ShowCursor;
sca
size(mousedata)
clf;
plot(mousedata(:,1), mousedata(:,2));
set(gca,'YDir','reverse');
axis equal
shg

```

Output 13.12.1a:

Trace the Circle clockwise, then click the left mouse button



This program is complicated, so it's best discussed piecemeal, which is why we have divided the program into parts using commented labels.

The first section, labeled *Initialization*, starts Psychtoolbox, sets some parameter values, and moves the mouse pointer to a specific x-y location on the screen with the function `SetMouse`. Next, a function called `HideCursor` is used to conceal the typical mouse cursor so it doesn't show up while your program is running.

The second part of the code, labeled *Mousewait*, makes the program pause until the mouse button is not depressed (which doesn't mean the mouse button is unhappy ☹). Without this code, the program might end immediately if the user happened to press the mouse button at the start of the program.


The third part of the code, labeled *Collect Data*, first sets a desired sampling rate, then starts a `while` loop that executes continuously until the user presses the left mouse button. This section of the code looks extremely complex, but you can work through the sequence of steps it performs. On each execution of the `while` loop, the following tasks are performed:

1. Draw the instructions and the large circle at a fixed position.
2. Draw the smaller cursor at the most recent position occupied by the mouse.
3. Flip the screen to show what has been drawn.
4. Record the new position of the mouse and whether mouse buttons are pressed, using `GetMouse`.
5. Store this new mouse position in the `mousedata` array.
6. Wait, using a short `while` loop, until it is time to get the next sample.

The reason the short `while` loop was used in the last step, instead (for example) of `Wait Secs(.1)` to sample 10 times a second, is that the timing is more precise. Each sample is taken exactly when it comes due. For an everyday example of the difference, suppose you wanted to check your e-mail once every hour. You would do so more precisely if you did so "on the hour" than if you checked and then waited for one hour before the next check. To see why this latter method is imprecise, consider that if it takes you 10 minutes to check your e-mail, you would actually check your email every 70 minutes instead of every 60 minutes using the second method.

When a mouse button click has been detected, the main `while` loop ends and the final part of the program, labeled *Cleanup*, is executed. First, the size of the `mousedata` array is reduced to the number of mouse samples collected. Next, the regular mouse cursor is reactivated with `ShowCursor`. At the end of the program, the variable `mousedata` contains a list of the mouse data points recorded during data collection. These data points are plotted to reconstruct the trajectory of the mouse. A new command, introduced here, `set(gca, 'YDir', 'reverse')`, inverts the direction of a specified axis. This command is used so the trace shown in the plot mirrors the direction of the mouse movements.

The *Cleanup* section also performs some arithmetic on the results. If all is well, the number of samples should agree with the sampling rate times the duration of the movement. The program reports both the desired sampling rate and the actual sampling rate.

Output 13.12.1b

```
DesiredSampleRate =  
    10  
ElapsedTime =  
    8.7009  
NumberOfSamples =  
    87  
ActualSampleRate =  
    9.9990
```

The above output shows that the sampling rate, 10 per second, was within the capability of the program. The actual sampling rate was almost exactly 10, but you can't assume that will always be the case.

Suppose you wanted to sample more frequently than 10 times per second. You could change the value of `DesiredSampleRate` from 10 to 100, say.

Code 13.12.2:

```
...  
% Part Three: Collect Data  
DesiredSampleRate = 100    % Number of samples per second  
clear sampletime;  
...
```

Here is the printed output of this experiment.

Output 13.12.2:

```
DesiredSampleRate =  
    100  
ElapsedTime =  
    4.5254
```

```

NumberOfSamples =
    271
ActualSampleRate =
    59.8843

```

Even though the program now calls for 100 samples per second, the actual rate is only 60 samples per second. Despite the instructions, the program could not keep up with the requested rate of 100 samples per second. Why not?

The rate at which the program could execute the mouse-sampling loop was limited by the screen refresh rate of the computer it ran on, and the program was tested on a monitor with a 60 Hz refresh rate. Because every `Flip` command was postponed until the next screen was refreshed to avoid flicker, the sampling rate could not exceed 60 per second. We would not have known about this problem had we not double checked the sampling rate at the end of the trial.

The take-home lesson is that you should do all you can to ensure that your program *actually* does what you think it does. Your computer doesn't necessarily do what you assume it is doing. It is important to build in checks to make sure that what is really happening on your computer is what you intend.

13.13 Creating an Animation With Moving Dots

The speed of Psychtoolbox allows you to do exciting things with visual displays. For example, the code below gets 1,000 dots to move continuously. The scientific purpose of such a display is to study visual sensitivity to dot-motion coherence, that is, the ability to tell which direction most of the dots are moving. The proportion of dots moving in the same direction can be varied. The output shows just one static image from the animation.

Code 13.13.1:

```

% Part One: Initialization
Screen('Preference', 'VisualDebugLevel', 1);
[window Scrnsize]= Screen('OpenWindow',0);
fliptime = Screen('GetFlipInterval', window);
centerX = Scrnsize(3)/2;
centerY = Scrnsize(4)/2;
coherence = .5;
numdots = 1000;
dotspeed = 2;
dotpos = zeros(2,numdots);
dotdir = zeros(1,numdots);
boxsize = 600;
trialduration = 3.0;
uniformdirection = randi(4); %1 = down, 2 = left,
%3 = up, 4 = right
for(dot = 1:numdots)


```

```
    dotpos(1,dot) = randi(boxsize); %Horizontal starting
% position
    dotpos(2,dot) = randi(boxsize); %Vertical starting
% position
    if(dot < ceil(numdots*coherence))
        dotdir(1,dot) = uniformdirection*pi/2;
    else
        dotdir(1,dot) = rand*2*pi;
    end
end

% Part Two: Animation
starttime = GetSecs;
timestamps = zeros(trialduration/fliptime,1);
counter = 0;
while(GetSecs-starttime < trialduration)
    for(dot = 1:numdots)
        dotpos(1,dot) = dotpos(1,dot) + ...
cos(dotdir(dot))*dotspeed;
        dotpos(2,dot) = dotpos(2,dot) + ...
sin(dotdir(dot))*dotspeed;

        if(dotpos(1,dot) > boxsize)
            dotpos(1,dot)= dotpos(1,dot)- boxsize;
        end
        if(dotpos(2,dot) > boxsize)
            dotpos(2,dot) = dotpos(2,dot)- boxsize;
        end
        if(dotpos(1,dot) < 1)
            dotpos(1,dot)= dotpos(1,dot)+ boxsize;
        end
        if(dotpos(2,dot) < 1)
            dotpos(2,dot)= dotpos(2,dot)+ boxsize;
        end
    end
    counter = counter + 1;
    Screen('Drawdots',window,dotpos,2,[255,0,0],[centerX-...
        boxsize/2, centerY-boxsize/2],1);
    timestamps(counter) = Screen('Flip',window);
end

% Part Three: User Response
sca
Response = input(['Which direction did most of the dots move?' ...
    '\n1 = down, 2 = left, 3 = up, 4 = right:']);
if(Response == uniformdirection)
    'You are correct'
else
    sprintf('The correct answer was %d',uniformdirection)
end
```

Output 13.13.1:

You can run the program yourself, trying to detect the direction in which most of the dots are moving, and responding with one of four numbers, as specified in the response prompt. Now try reducing the value of the coherence variable to determine how low this value can be before you lose the ability to detect the motion direction.

Code 13.13.1 uses a Screen operation called `DrawDots`, which draws the dots rapidly. In this case, if your computer is fast enough, 1,000 dots are drawn at the same rate as your monitor refresh rate (approximately every 17 milliseconds if your monitor's refresh rate is 60 Hz). To check what the animation speed actually is, you can look at the differences between subsequent values in the matrix `timestamps`, because these values are the times at which the `Flip` statements finished.

Here are more details about Code 13.13.1. In the first section, labeled *Initialization*, Psychtoolbox is started with `OpenWindow`, discussed earlier. In addition to the window pointer, a second variable called `Scrnsz` is also returned. This variable contains the display resolution of your monitor in pixels. The next two lines use these numbers to calculate the coordinates of the screen's center. After this, parameters are set to determine what proportion of the dots move in the same direction (the variable named `coherence`), the number of pixels that the dots move on every update of the display (`dotspeed`), the number of dots (`dotnum`), the size of the square area containing the dots (`boxsize`), the duration of the trial in seconds (`trialduration`), and the direction that the mass of coherent dots will move (`uniformdirection`). Each of the 1,000 dots is assigned a movement direction (`dotdir`) and an initial position (`dotpos`) specified in `x` and `y` coordinates.

The second part of the program, labeled *Animation*, specifies a `while` loop that runs for a predetermined time. On each execution of the `while` loop, four tasks are carried out. First, the position of each dot is updated according to its movement direction. Second, it is determined whether each dot has moved out of the square region. If so, that dot is made to "wrap around" to the other side of the square. Third, the dots are displayed using `DrawDots`. Fourth, the display is flipped and the timestamp of that flip is recorded. This final

part of the program shuts down Psychtoolbox, asks the subject to provide a response, and gives feedback about whether the response was correct.

13.14 Making Things Transparent

Color specifications in Psychtoolbox can include a fourth value that specifies a transparency level, called the *alpha* channel, which can be used to create stimuli that are fully or partially transparent. The following example illustrates how this works. We will not show the output here because the printed output would not do justice to the rendered graphics. However, you can run the program or see it on the text's website.

Code 13.14.1:

```
Screen('Preference', 'VisualDebugLevel', 1);
window = Screen('OpenWindow', 0, [150, 150, 150]);
Screen('Blendfunction', window, GL_SRC_ALPHA, ...
      GL_ONE_MINUS_SRC_ALPHA);
Screen('FillOval', window, [0, 0, 255, 75], [200, 200, 350, 350]);
Screen('FillOval', window, [0, 255, 0, 75], [300, 200, 450, 350]);
Screen('FillOval', window, [255, 0, 0, 75], [250, 250, 400, 400]);
Screen('Flip', window);
KbPressWait(-1)
sca
```

There is a new `Screen` operation here called `Blendfunction` that configures Psychtoolbox to use the alpha channel as transparency. At this point it is not important that you understand exactly what this command does. However, if you wish to learn more about other things that the alpha channel can do, look at the `Screen` documentation for `BlendFunction`. Be forewarned that it is not for the faint of heart!

Once transparency has been configured by `BlendFunction`, you can specify a fourth color value that specifies how transparent a shape will be, with 255 being fully opaque and 0 being fully transparent (i.e., invisible). In this example, the value is 75, which indicates about 30% transparency. Note that transparency also works with setting colors in `DrawText` and most other Psychtoolbox commands that specify colors. To learn more about transparency, adjust this value for one or more of the circles and run the code again. Remember that the order in which the circles are drawn influences how they appear when overlapping.

Code 13.14.1 has another new function called `KbPressWait`, which uses `KbCheck` to wait for a key press (with the argument of `-1` to accommodate any external keyboard). Check the help documentation for this function to learn more about it.

Textures afford an even more powerful capability. They let you control the transparency of each pixel in an image. This tool lets you make some parts of a texture more transparent than others. To see an example of this feature in action, run the following demonstration program that comes with Psychtoolbox. The output is not shown below.

Code 13.14.2:

```
| AlphaImageDemo
```

You can learn more about the demos that come with Psychtoolbox in Section 13.16.

13.15 Testing the Simon Effect With Psychtoolbox

Earlier in this book, in Code 12.6.5, we showed you how to use MATLAB to create a complete experiment to measure the Simon effect, the tendency for choice reaction times to be affected by aspects of a stimulus (typically spatial aspects) that are irrelevant to the stimulus' designation as a correct response. Code 12.6.5 used basic MATLAB commands to create stimuli and collect responses. What follows is a version of the same experiment using Psychtoolbox. The data file format is the same as in the previous version, so you should be able to analyze your data files using the same analysis program. The present version of the experiment introduces an additional requirement. The participant has just 2 seconds to respond. If no response occurs during that time, the program moves on to the next trial. This capacity to wait for a fixed time for a response is made possible through `KbCheck`, which checks the status of the keyboard without pausing the program.

Code 13.15.1:

```
% Part One: Initialization
Screen('Preference', 'VisualDebugLevel', 1);
sinit = input('Subject's initials: ','s');
outfilename = ['SimonDataPTB_' sinit];
[window Scrnsize]= Screen('OpenWindow',0);
halfFlip = Screen('GetFlipInterval', window)/2;
KbName('UnifyKeyNames');
centerX = Scrnsize(3)/2;
centerY = Scrnsize(4)/2;
Screen('TextFont',window, 'Arial');
Screen('TextSize',window, 72);
timeout = 2;
[ttype(1:4).side] = deal('L','R','L','R');
[ttype(1:4).stim] = deal('L','L','R','R');
[ttype(1:4).comp] = deal('C','I','I','C');
% Prepare data fields for each type of trial.
[ttype(1:4).RT] = deal([]);
[ttype(1:4).error] = deal(0);
%get the size of the fixation cross
[bounds] = Screen('TextBounds', window, '+');
fixSizeX = bounds(3)/2;
%get the size of the stimuli
[bounds] = Screen('TextBounds', window, 'L');
```

```
stimSizeX = bounds(3)/2;
leftStimX = centerX-400-stimSizeX;
rightStimX = centerX+400-stimSizeX;
ListenChar(2);

% Part Two: Data Collection
HideCursor
for blocknumber = 1:8
    for typenum = randperm(4);
        WaitSecs(2); %pause at start of trial, then show fixation
        Screen('DrawText', window , '+',centerX-fixSizeX ,...
            centerY,[0,0,0]);
        onsetTime1 = Screen('Flip',window);
        %Draw the fixation cross
        Screen('DrawText', window , '+',centerX-fixSizeX ,...
            centerY,[0,0,0]);
        %Show the stimulus on the left or right side
        if ttype(typenum).side == 'L'
            Screen('DrawText', window , ttype(typenum).stim,...
                leftStimX,centerY,[0,0,0]);
        else
            Screen('DrawText', window , ttype(typenum).stim,...
                rightStimX,centerY,[0,0,0]);
        end
        Starttime = Screen('Flip',window,onsetTime1 + 1.0 ...
            - halfFlip);
        Nowtime = Starttime;
        responseGiven = 0;
        response = 0;
        %collect a response with a timeout
        while(Nowtime < Starttime + timeout & responseGiven == 0)
            %Check for a response
            [keyDown secs keyCode] = KbCheck(-1);
            if(keyDown)
                responseGiven = 1;
                response = KbName(keyCode);
            end
            Nowtime = GetSecs; % check the current time
        end
        thisRT = secs-Starttime; %compute the reaction time

        if(response(1)=='a') %convert the response into L or R
            thisResp = 'L';
        elseif(response(1) == ';'')
            thisResp = 'R';
        else
            thisResp = 'X';
        end
    end
end
```

```

        if ttype(typenum).stim == thisResp
            ttype(typenum).RT = [ttype(typenum).RT thisRT];
        else
            ttype(typenum).error = ttype(typenum).error + 1;
            Beeper;
        end
        Screen('Flip',window);
    end
end

% Part Three: Cleanup and File save
ShowCursor
ListenChar(0);
sca
save(outfilename, 'ttype');

```

An example of the stimulus for an incompatible trial, with “L” to the right of the visual fixation cross, is shown here.

Output 13.15.1:

+ L

At this point, you should be able to study this program and figure out how it works. Comments have been added to help you. Being helped by them will remind you of how important it is to provide comments to make code more understandable.

The program uses two Psychtoolbox commands that were not introduced before. `TextBounds` gives the horizontal and vertical dimensions, in pixels, of a character string given the current font and font size. This information can be useful for tasks like centering a text stimulus at a location given that different fonts have different character widths.

Second, you will also note that we use the Psychtoolbox function `Beeper` instead of MATLAB’s `beep`. This is required because, at least on some platforms, Psychtoolbox is incompatible with MATLAB’s `beep` command.

13.16 Exploring Psychtoolbox Further

Psychtoolbox has many capabilities we didn’t touch on here. For example, using `Screen’s OpenMovie` command lets you show movie files while simultaneously collecting key

presses or mouse movements. With `PsychPortAudio`, you can play audio files with accurate timing and also record audio input using your computer's microphone.

Still other functions allow you to get your computer to interface with other devices for collecting data, such as eye-movement recording devices, EEG recording devices, game pads, and joysticks. You can find a list of such hardware interface programs and the devices for which they have been shown to work at <http://docs.psychtoolbox.org>.

Another useful aspect of Psychtoolbox is that an extensive collection of demonstration programs has been developed that illustrate the use of Psychtoolbox's more complex features. You can access the list of these demos as follows.

Code 13.16.1:

```
| help PsychDemos
```

The output will list all of the Psychtoolbox demos. You can run each one by typing its name in the Command window. The MATLAB code for each demo is also within MATLAB's path, so you can open the MATLAB files in the Editor. For example, try running `DriftDemo` and then open the code in the Editor with the following command:

Code 13.16.2:

```
| edit DriftDemo
```

Taking the time to explore these demos will help you learn about the impressive capabilities that Psychtoolbox affords.

13.17 Recovering From Psychtoolbox Program Crashes and Infinite Loops

It is anticlimactic to end this chapter on a “crashy” note, but we must do so. The reason is that, in Psychtoolbox, getting stuck in the middle of a still buggy program can be very problematic, not just because the program doesn't work but also because it interferes with your ability to interact with your computer. You will know you are in this unhappy state if your program becomes unresponsive or you hear a typical MATLAB error-beep from the computer but can't see the error message, and nothing you do to right the wrong has any apparent effect (or nothing is happening when something should be).

If you encounter such a situation in a standard MATLAB program, you can activate the command window to view the error message or interrupt the program with `ctrl-c`. If you encounter such an error while using Psychtoolbox, the first thing to do is, similarly, to return to the MATLAB Command window. However, because Psychtoolbox controls your screen, that window will be invisible. To bring it back, you need to take three steps. First, you need to make MATLAB the active window. On a Mac, press `command` and 0 (*zero*,

not *oh*) simultaneously to make the Command window the active window (although you will not actually see any change because the Psychtoolbox screen is still active). On a Windows computer, hold ALT and TAB down at the same time (and you may need to press Alt-Tab multiple times if you have multiple MATLAB windows open). On a Linux computer, press `ctrl-alt-esc`, followed by a mouse click. Note that doing this will not allow you to see the MATLAB window yet, which can be disconcerting. However, having made the MATLAB window active, you can now send commands to your computer even though you can't see them. The second step is to press `ctrl-c` at least three times: first, to interrupt any ongoing process; second, to ensure that `ListenChar` is not blocking keyboard input; and third, to clear any extraneous input in the command window. As a final step, type `sca` and press Return. This is short for `Screen('Close All')`. This series of inputs will tell MATLAB to close the Psychtoolbox screen, and you should find yourself comfortably back in the MATLAB environment.

In some rare cases, other programs or figure windows may interfere with your ability to exit Psychtoolbox using the steps just outlined. If this happens, stronger medicine may be needed. On Windows, you can press `ctrl-alt-delete` together to open the Windows task manager and then force MATLAB to exit. On a Macintosh, hold the command-option and escape keys down at the same time to force-quit MATLAB. On Linux, you may need to configure your own keyboard command to force-quit an application using the system preferences.

13.18 Problems

Problem 13.18.1:

The `DrawTexture` operation can place a texture multiple times. Modify the code of 13.7.1 to place three copies of the same image at different locations on the screen simultaneously at three different locations and at three different sizes. To resize the image, you will need to know its original dimensions which you can get from the `imagedata` variable before its cleared. You will need to use the parameters of `DrawTexture` named `sourceRect` and `destinationRect`, which you can read about in the documentation. Each of these is a vector containing four numbers that specify the corners of the image, just like the location parameter for `FrameOval` in Code 13.6.1. The variable `sourceRect` specifies the part of the texture you are copying from and the `destinationRect` specifies the destination you are copying to on the screen. If the destination rectangle is a different size than the source rectangle, the texture will be grown or shrunk automatically as appropriate so it fits.

Problem 13.18.2:

Modify your solution to 13.18.1 by adding a `for` loop so 10 copies of the texture appear on the screen in sequence, each one rotated by 36 degrees relative to the previous one and on the screen for precisely 300 milliseconds.

Problem 13.18.3:

Apparent motion occurs when a stimulus appears to move between two locations even though the stimulus is shown statically at one location, A, then statically at another location, B, then statically at A again, then statically at B again, and so on. Create a Psychtoolbox program to draw a circle that alternates between two locations repeatedly until you press any key, at which point the program exits. Make the circles 20 pixels in diameter and have the program wait 150 milliseconds between jumps. Use the method illustrated in Code 13.8.3 to ensure that your stimulus timing is precise. Modify the number of pixels between the two presentations of the circle to find the critical distance at which the circle appears to move back and forth, rather than blink on and off. You will then have created an apparent motion demonstration.

Problem 13.18.4:

Modify your solution to 13.18.3 so one presentation of the circle is presented at a fixed location and the other circle is presented at the current mouse position. This should allow you to control the separation of the circles with great precision. Use your program to find the largest separation at which two separate dots appear as one dot moving back and forth, by adjusting the separation using the mouse. Now, modify your program to use `DrawText` to display, near the bottom of the screen, the number of pixels between the centers of the two dots (calculated with the Pythagorean theorem). Use your program to discover the maximum separation distance, measured in screen pixels, at which the apparent motion illusion can occur. Remember to keep your eyes a fixed distance from the monitor while viewing the stimuli in the experimental display. Measure this distance because you will need it for the next problem.

Problem 13.18.5:

Use the pixel distance in 13.18.4 to compute the velocity at which the apparent motion illusion occurs. First, figure out the pixel density of your monitor by obtaining the pixel resolution of your monitor. You can use the command at the beginning of Code 13.13.1 to get the screen size if you don't already know it. Then measure the horizontal width of your screen using a ruler. Compute how many pixels there are in 1 centimeter of your screen. Using this value, and the pixel count from problem 13.18.4, convert your pixel count to centimeters. Because each jump occurs at 150 millisecond intervals, you can compute the velocity of the dot in terms of centimeters per second. The final step is to convert centimeters into degrees of visual angle, a typical unit of measurement in vision experiments. For this calculation you will also need to measure how far your eyes were from the monitor when you measured the threshold in 13.18.4. You can find tutorials and tools for this calculation with an internet search for 'visual angle.' What you should end up with is a measure of velocity in units of degrees of visual angle per second. This velocity is the threshold for apparent motion in this case. Below this velocity, you will observe motion. Above this velocity, the dots appear to blink on and off in stationary positions.

Problem 13.18.6:

Modify Code 13.10.1 so it reports the name of the key that was pressed and the reaction time relative to the start of the `while` loop.

Problem 13.18.7:

Modify the code of 13.12.1 so that the `while` loop ends when the user has completed the circle, rather than when the mouse button is clicked. To do this, you will need to draw a marker on the circle to remind the user of the point they started at. You can use several Psychtoolbox functions to do this, but we suggest `DrawLine`. You will also need to compute the distance between the mouse position and this starting point, and end the main `while` loop when that distance is sufficiently short.

Problem 13.18.8:

Create a program to determine the minimum duration that a stimulus has to be on the screen to be seen if it is followed by another stimulus. Your program should specify a list of four-letter words in a cell array. It should then pick one of the words randomly and display it in the center of the screen, in uppercase letters. After 100 milliseconds has elapsed, replace the word with the string `####` at exactly the same spatial position as the word, to serve as a mask. Leave this mask on the screen for exactly 1 second and then exit the program. Now, use this program and modify the 100 millisecond duration to determine the shortest duration at which you can still identify which one of the words was presented. Ensure that your program has accurate timing by using the techniques shown in Code 13.8.3.

Problem 13.18.9:

If your monitor has an adjustable refresh rate, change the refresh rate to be as fast as possible and then modify the Code 13.12.1 to determine the maximum rate of mouse sampling. Now deactivate all of the screen drawing commands as well as the `Flip` command so that nothing is drawn on the screen. What is the maximum rate of mouse sampling in this case?

Problem 13.18.10:

Modify the code to 13.15.1 to provide the user with auditory feedback using `PsychPortAudio` instead of `Beeper`. To figure out how to use `PsychPortAudio`, consult the documentation and the demonstration program that comes with Psychtoolbox: `BasicSoundOutputDemo`. There is also a helpful Psychtoolbox function named `MakeBeep` to do some math for you. Make a beep with a frequency of about 800 Hz and that lasts for about 300 milliseconds.

Problem 13.18.11:

Modify your answer to 13.18.10 so instead of playing a beep, the computer says “wrong” through its speaker. You will need a working microphone hooked up to your computer for this one. Use the `BasicSoundInputDemo` from Psychtoolbox to record a sound sample of someone saying “wrong,” which will be saved as a `.wav` file. You will then need to load this file at the top of your experiment using MATLAB’s `wavread` function, and send the resultant audio data to the audio buffer, as you did in problem 13.18.9. (Mac OS users can try the command `!say wrong` as an alternative to using `PsychPortAudio`.)

14. Debugging

This chapter covers the following topics:

- 14.1 Debugging using error messages and breakpoints
- 14.2 Using temporary feedback for debugging
- 14.3 Interpreting error messages
- 14.4 Using graphic output for programming and debugging
- 14.5 Practicing debugging

The commands that are introduced and the sections in which they are premiered are:

<code>dbclear</code>	(14.1)
<code>dbcont</code>	(14.1)
<code>dbquit</code>	(14.1)
<code>dbstep</code>	(14.1)

14.1 Debugging Using Error Messages and Breakpoints

In an ideal world, every program you write would be perfect from the moment your fingers touch the keys. Every character you type would be exactly right. Colleagues peering over your shoulder would marvel at the speed with which you go from an initial idea, conceived in an instant, to a MATLAB masterpiece entered with virtuosity at the keyboard.

Ah, the fantasy! The truth is that just as writing is rewriting—a well-known mantra of authors—programming is “reprogramming.” Program development in real life is a cyclic process of writing code, thinking or hoping it’s right, then getting your wrist slapped, and then rewriting the code, and going through this process over and over. It bugs programmers that they’re imperfect, but all of them know, and the authors of this book certainly know, that in real life, programming involves debugging. Given how central debugging is for (MATLAB) programming, we have written an entire chapter about this process.

Why have we put this chapter near the end of the book rather than near the beginning? Our rationale is that working through the examples we want to convey here depends on familiarity with the MATLAB commands used in them. If you have worked through the book to this point, you have already done a great deal of debugging. The purpose of this chapter is to point out some techniques that you may not have discovered, invented, or learned from others.

Here is an example of a program with problems typical of the first draft of a MATLAB program. The goal is to make an array, `a(1:6)`, of the squares of the first six integers, then another array, `b(1:6)`, of their square roots, and finally, report the values of `a` and `b`. The initial program has a couple of errors, and several other errors come to light in the process of tracking them down. We use this as an example of the iterative nature of debugging. Before reading further, you might examine Code 14.1.1 and see if you can see what

the problems might be. Type it in or download it from the book's website. Then, try to make the program work as intended, before reading further in the description.

Code 14.1.1:

```
% Code_14_1_1
function main
a = [1:6]^2
makeb;
a;
b;
end

function makeb
for a = 1:6
    b(a) = sqrt(a)
end
```

As it stands, the program fails.

Output 14.1.1:

```
Error: File: debug1.m Line: 14 Column: 1
The function "main" was closed with an 'end', but at
least one other
function definition was not. To avoid confusion when using
nested
functions, it is illegal to use both conventions in the
same file.
```

Following the hints of the error message, we add another `end` to the `makeb` function. (The `end` that is already there goes with the `for` statement, not the function statement.) As we type in the `end`, the function statement of `makeb` highlights briefly, reassuring us that we've put the `end` in the right place. We try again.

Code 14.1.2:

```
% Code_14_1_2
function main
a = [1:6]^2
makeb;
a;
b;
end

function makeb
for a = 1:6
    b(a) = sqrt(a)
end
end
```

Output 14.1.2:

```
Error using mpower
Inputs must be a scalar and a square matrix.
To compute elementwise POWER, use POWER (.^) instead.
```

We now add one character to the second line and try again.

Code 14.1.3:

```
% Code_14_1_3
function main
a = [1:6].^2
makeb;
a;
b;
end

function makeb
for a = 1:6
    b(a) = sqrt(a)
end
end
```

Output 14.1.3:

```
a =
     1     4     9    16    25    36
b =
     1
b =
     1.0000     1.4142
b =
     1.0000     1.4142     1.7321
b =
     1.0000     1.4142     1.7321     2.0000
b =
     1.0000     1.4142     1.7321     2.0000     2.2361
b =
     1.0000     1.4142     1.7321     2.0000     2.2361     2.4495

Undefined function or variable 'b'.
Error in debug2 (line 6)
b;
```

Now what? A new problem has arisen! The variable `b` seems to be undefined when we return from the function, but the function has clearly defined it, as indicated by the fact that `b` prints out each time through the `for` loop in `makeb`. Can you spot the problem?

Because `makeb` is a local function, the definition of `b` within `makeb` is not visible to the main function. `makeb` needs to be a nested function, not a local one, so the end of `main` that we added in Code 14.1.2 belongs *after* `makeb`, not before. That way, `makeb` will be a nested function, and the main function can see its variables (and vice versa). We'll put a comment after the moved `end` to remind us of the function it is related to because it comes in a series of three ends in a row, which is potentially confusing. After we have invoked automatic formatting (see Section 2.6), the indentation of the `makeb` function reminds us that it is now a nested function.

Code 14.1.4:

```
% Code_14_1_4
function main
a = [1:6].^2
makeb;
a;
b;

    function makeb
        for a = 1:6
            b(a) = sqrt(a)
        end
    end

end %function main
```

Output 14.1.4:

```
a =
    1     4     9    16    25    36
b =
    1
b =
    1.0000    1.4142
b =
    1.0000    1.4142    1.7321
b =
    1.0000    1.4142    1.7321    2.0000
b =
    1.0000    1.4142    1.7321    2.0000    2.2361
b =
    1.0000    1.4142    1.7321    2.0000    2.2361    2.4495
```

The output looks much better, but we don't need to see `b` every time through the loop. In restoring the semi-colon to the statement that assigns `b`, we see that we also omitted a semi-colon in the statement that originally generated `a`, so we add a semi-colon in both places.

Code 14.1.5:

```
% Code_14_1_5
function main
a = [1:6].^2;
makeb;
a;
b;

    function makeb
        for a = 1:6
            b(a) = sqrt(a)
        end
    end

end %function main
```

Output 14.1.5:

```
| >>
```

Now we have good news but also bad news. There is no error messages this time, which is good, but there is no output either, which is bad. Can you see the solution?

We had originally put semi-colons on the lines that were to put out `a` and `b` (lines 5 and 6 of Code 14.1.5), so no output was generated. They are easily removed, so we do so.

Code 14.1.6:

```
% Code_14_1_6
function main
a = [1:6].^2;
makeb;
a
b

    function makeb
        for a = 1:6
            b(a) = sqrt(a)
        end
    end

end %function main
```


Output 14.1.6:

```
| a =
      6
b =
    1.0000    1.4142    1.7321    2.0000    2.2361
    2.4495
```

Much better, but where `a` had the value `[1 4 9 16 25 36]` earlier, it is now just 6. So *now* what's going on?

Line 5 should have generated the same result as in Code 14.1.4. To learn why it did not, we'll use a MATLAB debugging feature we haven't used before, a *breakpoint*. A breakpoint is a signal to MATLAB to run the program in the Editor window up to a particular line, and then stop just before executing that line. We'll click on the dash just to the right of the line number "4" in the left margin of the Editor window. Now we have a breakpoint, which shows as a little stop sign in the left margin. When the code is run, it stops just before executing line number 4 (`makeb`), with a green arrow pointing at that line to indicate where it stopped. In the Command window, the prompt `K>>` indicates that we have stopped the program in the middle of its execution.

Output 14.1.7:



While the code is stopped, we can examine or change the values of variables before we continue.

First, let's examine `a`, and print its current value:

Code 14.1.8:

```

K>> a
a =
     1     4     9    16    25    36

```

Next, we take one step forward in the program using `dbstep`, which causes line 4 to execute, and stops us at line 5, where the green arrow now points. We see all the output generated within the `makeb` function, followed by `5 a`, the line number and next command to be executed.

Code 14.1.9:

```

K>> dbstep
b =
    1
b =
    1.0000    1.4142
b =
    1.0000    1.4142    1.7321
b =
    1.0000    1.4142    1.7321    2.0000
b =
    1.0000    1.4142    1.7321    2.0000    2.2361
b =
    1.0000    1.4142    1.7321    2.0000    2.2361
2.4495
5   a

```

Another `dbstep` executes line 5, `a`, printing the value of `a`, followed by 6 `b`, the line number and next command to be executed, which is where the green arrow points next.

Code 14.1.10:

```

K>> dbstep
a =
    6
6   b

```

We now see that the value of `a` has been changed by the running of the `makeb` function. Knowing that this has happened can help us find the remaining problem(s) in the program. We'll leave it to you to find and fix them.

One final application of using a breakpoint and `dbstep` is to understand a working program. For example, you can best understand the sequence of operations of a recursive program similar to Code 8.7.3 by stepping through it one command at a time. Start by putting a breakpoint on the first command line. You can then use `dbstep` in the Command window (or the “step” button, which appears when a breakpoint is active) to trace the program's program flow from beginning to end, as indicated by the location of the green arrow in the left margin.

Now that you're done debugging, the `dbquit` command exits debugging mode. When that's done, you can remove the breakpoint by clicking on it.

Code 14.1.11:

```

K>> dbquit
>>

```

There's much more to learn about breakpoints in MATLAB, but this extended protocol analysis suggests some of the tools MATLAB provides for low-level debugging. You can consult the documentation in MATLAB and on the website to learn more about `dbcont` (run up to the next breakpoint), `dbclear` (clear all breakpoints), and other commands as you develop facility with the basic features.

There are other lessons to learn from this example. One is that virtually all programs have bugs initially. Virtually all programmers spend as much testing and fixing bugs than they do generating original code. Therefore, do not think, if you are a student, that you are in any way below par if you spend a lot of time debugging. You are doing what all programmers do. Second, bugs may have side effects. Fixing a bug in one place may cause or reveal a logical error elsewhere. Third, there may be *many* unrelated bugs in a program. You have to track them all down before your program can be relied on. Fourth and finally, debugging is an empirical process. Think of it as an experiment on the program you are working on to understand how the program actually works.

14.2 Using Temporary Feedback for Debugging

In the program we just debugged, if we had anticipated the problems we encountered, we might have chosen to generate output at several points for test purposes, planning to eliminate that feedback once we were sure the program ran correctly. Here we illustrate that approach.

In the code below, the Boolean variable `testing` controls whether or not there is intermediate output. The output can be suppressed by changing the assignment of `testing` in line 2 from `true` to `false`. We put all the temporary test commands into one line. This approach is clerically easier and less error-prone than selectively removing and replacing semicolons to control output during testing, or manually deleting printing commands.

Code 14.2.1:

```
function main
testing = true;
a = [1:6].^2;
if testing, disp('testing a'), disp(a), end;
makeb;
if testing, disp('testing a again'), disp(a), end;
a
b

    function makeb
        for i = 1:6
            b(i) = sqrt(i);
        end
    end

end %function main
```

Output 14.2.1:

```

testing a
      1      4      9     16     25     36
testing a again
      1      4      9     16     25     36
a =
      1      4      9     16     25     36
b =
      1.0000      1.4142      1.7321      2.0000      2.2361      2.4495

```

Once the program's accuracy is confirmed, it is easy to change the second line to `testing = false`, disabling any output lines that were shown during program development.

14.3 Interpreting Error Messages

Sometimes the error message to a file does not point directly to the problem. Here is code for a very simple program saved as `Code_14_3_1.m`. If you run it, you get pretty horrifying feedback from MATLAB, making it sound like you may have done major damage to your computer. Changing the recursion limit is not necessary, we assure you, nor would it solve the problem. The problem can be fixed by inserting a single character at a well-placed position. The take-home lesson is that sometimes error messages don't point to problem origins but instead point to problem consequences. This can be challenging. It takes experience to know what error messages mean in the contexts where they arise.

Code 14.3.1:

```

Code_14_3_1
x = 1

```

Output 14.3.1:

```

Maximum recursion limit of 500 reached. Use
set(0, 'RecursionLimit', N) to change the limit. Be aware
that
exceeding your available stack space can crash MATLAB
and/or
your computer.
Error in Code_14_3_1

```

In case it's not obvious, all that needs to be done is to turn the first line into a comment by typing a `%` sign before it or by clicking on it and hitting `ctrl-r` (command-`/` on the Mac).

14.4 Using Graphic Output for Programming and Debugging

Graphics are useful for more than just displaying final results. They are also useful for checking the accuracy of computations. Code 14.4.1 illustrates this approach in connection with an algorithm for detecting the duration of a transient spike in an analog signal, such as one that might be obtained in a single-unit neural recording study. Step 1 reads in the data and displays it, which is usually a good idea to make sure your data is reasonable. Step 2 determines the neighborhood of the spike by detecting the data's excursion through half its maximum. Step 3 repeatedly moves an index variable, *i*, to the left (starting from the abscissa value of the first excursion through half the spike amplitude) one step at a time, as long as the values of the spike are each smaller than the *following* one. In this way, the program identifies the point at which the spike begins to increase monotonically. Step 4 moves *i* to the right, starting from the last excursion as long as the values of the spike are each smaller than the *preceding* one, to identify the point at which the sample stops monotonically decreasing. Finally, Step 5 reports the spike duration, as determined by the interval between the first of the monotonically increasing points and the last of the monotonically decreasing points.

Code 14.4.1:

```
% Code 14_4_1
testing = true;

% 1. Read in and show the data
clc;
load('spikedata');
if testing
    figure(1); clf;
plot(xvals, 'k'); hold on
end

% 2. Detect spike half/amplitude excursion
GreaterThanHalf = xvals > max(xvals)/2;
plot(GreaterThanHalf, 'k-.');
peakvals = find(GreaterThanHalf);

% 3. Move to the left from first excursion
% as long as spike monotonically declines
firstval = peakvals(1);
while xvals(firstval) > xvals(firstval-1)
    firstval = firstval - 1;
    if testing
        plot(firstval, xvals(firstval), 'ko');
        pause(0.5)
    end
end

% 4. Move to the right from last excursion
% as long as spike monotonically declines
```


```

lastval = peakvals(end);
while xvals(lastval) > xvals(lastval+1)
    lastval = lastval + 1;
    if testing
        plot(lastval,xvals(lastval),'ko');
        pause(0.5)
    end
end

% 5. report the results
if testing
    plot([firstval,firstval],[-1,xvals(firstval)],'k--');
    plot([lastval,lastval],[-1,xvals(lastval)],'k--');
    text(140,5,sprintf(...
        ['The spike begins at sample %d \n' ...
        'and ends at sample %d.'], ...
        firstval,lastval),'fontsize',16);
    text(140,4,sprintf(...
        'Duration is %d samples.',...
        lastval-firstval),'fontsize',16);
    saveas(1,'Output_14_4_1.eps')
end

```

Output 14.4.1:




The succession of circles along the rising and trailing edges of the spike reassure us that the analysis is doing what we intend. This particular signal is not noticeably noisy. Suppose the data were noisy, however, in which case the algorithm might not work. The graphic could alert us to this shortcoming before we rush to publish. We can make the data used in Code 14.4.1 just a bit noisy by adding some randomness to the signal, but otherwise analyze it the same way.

Code 14.4.2:

```
% Code 14_4_2
testing = true;

% 1. Read in and show the data
clc;
load('spikedata');
xvals = xvals + randn(301,1)*.4;

% . . . the rest of Code 14.4.2 is unchanged from Code 14.4.1
```

Output 14.4.2:

Output 14.4.2 demonstrates that the algorithm underestimates the duration of the spike in the noisy data because it does not find the beginning of the spike and prematurely detects the end of the spike, due to the noise.

What to do in such a situation is up to your good judgment, imagination, and creativity. You could filter the data to attenuate the noise, and/or you could use a different definition for spike duration. For example, you could define excursions through 1/10 of the spike amplitude as indicating its beginning or end. Whatever algorithm you use, even if it is one that you have adapted from published research, you will profit from closely observing its operation while you develop your program, applying it to samples of data across both sessions and subjects to check that it is working as intended.

14.5 Practicing Debugging

Problem 14.5.1:

This program was designed to generate a million sums and measure how long it takes to do so. To test it, type it in exactly as printed here or get it from the website and paste it into your Editor window.

Code 14.5.1:

```
% Problem_14.5.1.m
clear;
m = zeros(100,100,100);
tic;
for i = 1:100
    for j = 1,100
        for k = 1:100
            m(i,j,k) = i + j + k;
        end;
    end;
end;
toc;
```

There is a problem with the program as written, which will become evident in the output when it is corrected. How does fixing the problem affect the program's operation? What "defensive programming" might be used to guard against the disastrous effects of such a slip of the finger? (Hint: There is a NaN-obvious solution.)

When the corrected program is before you, experiment with lines 2 and 3 to explore the effects on execution time of clearing or not clearing variables and pre-allocating or not pre-allocating memory for the variables.

Problem 14.5.2:

You write a program to analyze students' scores in a test. You test your program with a small set of scores, just four tests for each of 10 students. You are interested in the mean scores for each test for all of the students whose overall mean scores equal or exceed the grand mean of all the scores and, separately, the mean scores for each test for all of the students whose overall mean scores fall below the grand mean of all the scores. Your program appears below, along with the output you receive. You feel very proud of what you've done because, as expected, the students in the first group have higher mean test scores than do the students in the second group. However, your professor looks over your shoulder and shakes her head. "Whoops," she says. "Are you sure you got it right? Try removing the semi-colon after `ok_students` (before the final `end` statement)," she continues. "Maybe you could move a couple of lines of code." What did she mean? Revise the program and rerun it. In the design of the original program, what precaution might you have taken to ensure the problem would come to your attention before you submitted the solution to your professor?

Code 14.5.2:

```
clear all
clc
commandwindow

scores=[
    92 87 65 43
```

```
    86 86 71 22
    67 55 78 80
    70 65 58 98
    99 95 98 93
    88 80 72 90
    82 80 77 71
    90 90 89 90
    45 40 51 29
    77 77 78 81
  ]
  sz_scores=size(scores);
  ok_scores=[];
  ok_students=[];
  for pass=1:2
    for r=1:sz_scores(1)
      if pass==1
        if mean(scores(r,:))>= mean(mean(scores))
          ok_students=[ok_students r];
          ok_scores=[ok_scores;scores(r,:)];
        end
      else
        if mean(scores(r,:))< mean(mean(scores))
          ok_students=[ok_students r];
          ok_scores=[ok_scores;scores(r,:)];
        end
      end
    end
  end
  pass
  mean(ok_scores)
  ok_students;
end
```

Output 14.5.2:

```
scores =

    92    87    65    43
    86    86    71    22
    67    55    78    80
    70    65    58    98
    99    95    98    93
    88    80    72    90
    82    80    77    71
    90    90    89    90
    45    40    51    29
    77    77    78    81

pass =

    1
```

```

ans =
    87.2    84.4    82.8    85
pass =
     2
ans =
    79.6    75.5    73.7    69.7

```

Problem 14.5.3:

The following code is based on Code 3.8.4, but differs in an important respect. As in Code 3.8.4, a 1×4 matrix is expected. Make your prediction of the results, then check your prediction by executing the code. Hint: When you type the code into your Editor window, use copy–paste to repeatedly enter the variable name `matrix_to_be_appended_to` so you don’t have to type it in each time. That will help you avoid typos that need debugging.

```

matrix_to_be_appended_to = []
matrix_to_be_appended_to = [matrix_to_be_appended_to + 1]
matrix_to_be_appended_to = [matrix_to_be_appended_to + 2]
matrix_to_be_appended_to = [matrix_to_be_appended_to + 3]
matrix_to_be_appended_to = [matrix_to_be_appended_to + 4]

```

If the program does not work as you expected, experiment in the Command window to learn how to fix it so it does.

Problem 14.5.4:

There’s a problem in this code. Find it by using a breakpoint to stop just before executing the offending line, so you can use the Command window to figure out what the problem is. No fair just using your insight! Hint: stopping the program at just the right place using the breakpoint function will help you “size” up the problem.

Code 14.5.4:

```

a = zeros(10,100);
b = ones(10,100);
c = randi(10,100,9);
d = a + b + c;

```

15. Going On

This chapter covers the following topics:

- 15.1 Programming productively
- 15.2 Finding and navigating in the Editor
- 15.3 Double commenting
- 15.4 Comparing files
- 15.5 Profiling for efficiency
- 15.6 Examining built-in functions
- 15.7 Creating stand-alone applications
- 15.8 Programming ethically
- 15.9 Reading further

The commands that are introduced and the sections in which they are premiered are as follows:

`profile` (15.1)

`%% (section header)` (15.3)

15.1 Programming Productively

A lot of material has been covered in this book, and though you are about to “graduate,” it may be better to speak of “commencement” rather than “completion” at this time. We want to help you go on from here, capitalizing on what you have learned to make good decisions, and also wise (ethical) decisions, related to MATLAB programming. This chapter is designed to serve those purposes.

The first general topic covered here concerns programming productively. As you continue to work with MATLAB, you will discover timesaving habits that will be useful to you as you generate bigger and more complex programs.

One piece of advice about programming productively is to find a programming style that works well for you. A book by Johnson (2011) offers helpful suggestions about MATLAB programming style. There is no one best style for everyone, however. As you have seen here, programs can take different forms depending on the particular needs they address and also, as it happens, depending on who writes the program. Some of those stylistic differences have been reflected in this book.

Besides having a style that you prefer, you should cultivate tools that can facilitate your programming. The next sections cover some of these. Others can be found by exploring MATLAB’s menus, by using MATLAB’s `help`, by reading MATLAB’s docs (accessible via the `doc` command), by turning to the MathWorks’ web pages, and, perhaps most importantly, by interacting with others who program.

15.2 Finding and Navigating in the Editor

Here are some helpful editing hints, in no particular order:

A time sink you can avoid is hunting laboriously for segments of code to be changed. When you are looking for a particular segment of code in a long program, you can use the Find button in the Editor. This can save time and eyestrain.

To change the name of a variable everywhere it occurs—for example, to make it more meaningful—you can use the Find & Replace window of the Find button to make the change all through the program.

Less obviously, if you want to look for code you believe you wrote or saw in one or more saved files, you can use the Find File button in the Editor. There, you can use the “*find files containing text*” option to find all the instances of that piece of code in all the MATLAB files of the current folder, if the current folder is the domain of the search. There are other options, however. To look for all .m files that begin with the same string, such as 'Lanyun' (the graduate student with whom the first author was doing quite a bit of programming at the time of this writing, albeit on a different project), type Lanyun*.m in the “*find files named:*” box.

If you have jumped to a remote section of your program to make a change, there is a back arrow button in the Editor toolbar that will return you to your point of departure.

If you are working on two distant parts of the program at once, you can split the editor screen horizontally to see both parts of your code simultaneously.

15.3 Double Commenting

If you want to mark an important location in your program, such as the beginning of a nested or local function that you may want to easily find and come back to, use the “two percent” solution. The %% comment has a special function in MATLAB. The commented line stands out because it is automatically emphasized in bold face, and it defines a *section*. It is easy to navigate to the beginning of a section via a button at the top of the Editor window that lists all the section headings in the program. Sections have other features that are useful to sophisticated programmers. We won’t go into them here, but to learn more, search for “MATLAB Run Code Sections” on the Internet. Here is how %% comments might have been used in part of Code 12.5.5.

Code 15.3.1:

```
function SimonDemo;
clc
clear
close all;
%% File Setup
sinit = input('Subject's initials: ','s');
```

```
outfilename = ['SimonData_' sinit];
rawdataoutfilename = strrep(outfilename, '_', '_Rawdata_');
rawdataoutfilename = strcat(rawdataoutfilename, '.txt');
rawdatafile = fopen(rawdataoutfilename, 'w');
fprintf(rawdatafile, ...
        'Trial\tside\tstim\tcomp\tKey\tResp.\tRT\n');
%% Setting the Window
screensize = get(0, 'screensize');
hfig = figure(...
    'position', [0 0 screensize(3) 200], 'color', [1 1 1]);
...
```

15.4 Comparing Files

If you have modified a previous program and the new version does not work, or you just want to know how it differs from the old version, you can use the Compare button of the Editor. The two files will be listed side by side, highlighting every difference between them (added or deleted lines, as well as changes within lines).

15.5 Profiling for Efficiency

If you have a program that takes a very long time to run, MATLAB provides a function called `profile` that lets you determine how long the components of your programs take to execute as well as other potentially useful information about your program. This function can be useful when you want to find out where your program is spending most of its running time. See MATLAB Help for more information about `profile`.

15.6 Examining Built-In Functions

Another thing to keep in mind is that you can open and read many of the functions (the built-in `.m` files) provided by The MathWorks. Sometimes it is helpful to do this so you can inspect these functions and see how the “maestros” at The MathWorks designed the functions. There may be times when you’d like to make a copy of such a function and edit it for your own needs. If you edit any MathWorks-supplied function, we strongly recommend that immediately after opening the file, you save it with a new name to ensure that you leave the original function untouched. For example, save `max` as `my_max` if you feel that you must modify the MathWorks-supplied `max` function. We recommend saving such a personalized copy of a built-in function even if you only intend to *read* the function. Accidentally modifying it in a way that makes it dysfunctional can cause you lots of grief.

15.7 Creating Stand-Alone Applications

You can write MATLAB programs that can be run as stand-alone applications to be run on computers that do not have, or by people (or computer accounts) who do not have,

MATLAB. To do this, you need the MATLAB Compiler toolbox. See the MathWorks website (www.mathworks.com/) for more information.

15.8 Programming Ethically

This next-to-last section of this chapter covers a topic that is rarely mentioned in computer programming textbooks, but it is one we feel strongly about, so we devote a fair amount of space to it.

In this book, we provided you with a great deal of technical information about how to program in MATLAB. We had other aims as well. One was to help you hone your thinking skills. As you have seen, when you program, you must be explicit. The “creature” you are dealing with, the computer, knows nothing about you or your intentions. The computer takes every line of code you write and cuts you no slack for the kind of day you’ve had, whether you donated to the poor, and whether, through your research, you are trying to solve a practical problem on which many lives depend. If you violate some rule of MATLAB syntax, you will get the same error message regardless of whether you are a saint or a scoundrel.

Why say this? The reason is that with the skill you have hopefully acquired here, you now have the power to do whatever you want, computationally speaking. But you can also, given your newfound knowledge, pursue considerable good or evil. If you wanted to—and of course we hope you won’t—you could wreak havoc through MATLAB. By drawing on your knowledge of this programming language, coupled with your knowledge of statistics, you could, if you were so inclined, make up data whole cloth. You were exposed to this practice in this book. We showed you hypothetical data used to illustrate programming techniques in many places. The aim of the simulations was to see whether putative processes and their associated parameters (e.g., presumed rates of memory decay) corresponded to *real* data. This is actually a time-honored way to evaluate theoretical models, provided the fabricated nature of the data is made explicit. For a review of modeling, see Busmeyer and Diedrich (2010) and Lewandowsky and Farrell (2011).

If you have less than honorable intentions, you could, as we just said, use MATLAB to make up data whole cloth. You could do this with virtual impunity by generating pseudo-data that are convincing by virtue of their resemblance to actual results. Your fake data could exhibit means that fall within reasonable bounds, express plausible patterns of main effects and interactions, exhibit typical patterns of variability, and, in general, could be assembled in a way that avoids the specter of being “too good to be true.”

Data that are too good to be true have alerted sharp-eyed investigators to their falsity. The most famous example in behavioral science was the data set of Cyril Burt, the British educational psychologist who claimed, via a supposed study of large numbers of identical twins separated at birth, that their behavioral similarities were too great to be due to nurture. “Nature, not nurture, accounted for variations in intelligence,” Burt declared (or words to that effect—this is not a direct quote). However, Leon Kamin (1974) of Princeton University spotted features of Burt’s data that made him suspicious of their veracity. Ultimately, Kamin showed that some of Burt’s data were fabricated. Had Burt known (or had access) to MATLAB or some analogous program, he might have escaped the notice of sleuths like Kamin. The same could be said for other researchers who, subsequently,

were found out because features of their data gave away their data's sordid origins. Among these researchers was Gregor Mendel, the father of genetics, whose data were shown to be implausibly perfect by Ronald Fisher (1936), the father of statistics, after whom the *F* statistic was named. We obtained the reference to Fisher (1936) from a Wikipedia article about Gregor Mendel.

In some cases, fraud in science has been detected because of irregularities in the way data were collected. This was what happened in the two most famous recent cases of fraud in behavioral science—the case of Marc Hauser, formerly of Harvard University, and the case of Diederik Stapel, formerly of Tilburg University. Articles about both of these now disgraced individuals can be found in Wikipedia. In both cases, coworkers became suspicious of the astonishing productivity of the scientists because, among other things, the scientists published far more data than could be vouched for.

You may now know enough about MATLAB to become unbelievably productive yourself. You might even be clever enough to temper your productivity so your data are not only not too plentiful to be believed but also imperfect enough to seem real.

We say these things not to “give you ideas,” nor to put the “fear of God” in you, but instead to encourage you to use your newfound powers for good. To the extent you may be susceptible to temptation, be aware of the fact that computational tools have recently been developed for flagging suspicious data (Enserink, 2012; Simonsohn, 2013). This an example of programming for good rather than evil.

We hope you will use your programming skills for good purposes as well. That you should is a reflection of the fact that you should, as a matter of course, choose good over bad. But leaving aside whatever “good” is and whatever “bad” is, we think we can say that you will be safely guided by an attitude that has underlain virtually page of this book: Humility is a virtue.

As we have indicated here, nothing is quite so humbling as believing as you have written flawless code only to discover that it has mistakes. Getting feedback from the computer that you are imperfect can reinforce your modesty. By extension, whatever you may believe about the supreme correctness of your understanding of behavioral science, there is some chance the hypotheses you dream up may actually be wrong. Don't feel, then, that you are above the law (behavioral or otherwise). Instead, via the humble act of programming, simply do your best, as honestly as you can, to contribute as best you can.

15.9 Reading Further

A number of other sources can be used to supplement the material covered in this book. They are listed in the References that follow and in the other sources that have been mentioned here.

References

- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Busmeyer, J. R., & Diedrich, A. (2010). *Cognitive Modeling*. Los Angeles, CA: Sage.
- Cohen, R. G., & Rosenbaum, D. A. (2004). Where objects are grasped reveals how grasps are planned: Generation and recall of motor plans. *Experimental Brain Research*, 157, 486–495.
- Dweck, C. S., & Bempechat, J. (1980). Children's theories of intelligence: Consequences for learning. In S. G. Paris, G. M. Olson, & H. W. Stevenson (Eds.), *Learning and Motivation in the Classroom*. Hillsdale, NJ: Erlbaum.
- Elliott, M. T., Welchman, A. E., & Wing, A. M. (2009). MatTAP: A MATLAB toolbox for the control and analysis of movement synchronisation experiments. *Neuroscience Methods*, 177, 250–257.
- Enserink, M. (2012, July 6). Fraud-detection tool could shake up psychology. *Science*, 337, 21–22.
- Fine, I., & Boynton, G. (2013). *MATLAB for the Behavioral Sciences* [Kindle ed.]. www.amazon.com/Matlab-Behavioral-Sciences-ebook/dp/B00CPT86NC
- Fisher, R. A. (1936). Has Mendel's work been rediscovered? *Annals of Science*, 1, 115–137.
- Hayes, B. (2012). Murkiness in numerical computing. *American Scientist*, 100(1), 84. doi: 10.1511/2012.94.84.
- James, W. (1890). *Principles of Psychology*. New York, NY: Holt.
- Johnson, R. K. (2011). *The Elements of MATLAB Style*. New York, NY: Cambridge University Press.
- Julesz, B. (1971). *Foundations of Cyclopean Perception*. Chicago, IL: University of Chicago Press.
- Kamin, L. J. (1974). *The Science and Politics of IQ*. Potomac, MD: Erlbaum.
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? [ECVP Abstract Supplement 14]. *Perception*, 36.
- Lewandowsky, S., & Farrell, S. (2011). *Computational Modeling in Cognition: Principles and Practice*. Los Angeles, CA: Sage.
- Lu, C.-H., & Proctor, R. (1995). The influence of irrelevant location information on performance: A review of the Simon and spatial Stroop effects. *Psychonomic Bulletin and Review*, 2, 174–207.
- MacLeod, C. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, 109, 163–203.
- Madan, C. (2014). *An Introduction to MATLAB for Behavioral Researchers*. Thousand Oaks, CA: Sage.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Plant, R. R., & Quinlan, P. T. (2013). Could millisecond timing errors in commonly used equipment be a cause of replication failure in some neuroscience studies? *Cognitive, Affective, and Behavioral Neuroscience*, 13, 598–614. doi:10.3758/s13415-013-0166-6.
- Plant, R. R., & Turner, G. (2009). Millisecond precision psychological research in a world of commodity computers: New hardware, new problems? *Behavior Research Methods*, 41, 598–614. doi:10.3758/BRM.41.3.598
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (2007). *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press. (See also www.nr.com).
- Rand Corporation. (1955). *A Million Random Digits With 100,000 Normal Deviates*. www.rand.org/pubs/monograph_reports/MR1418.html
- Rock, I. (1985). *Perception*. New York, NY: Scientific American Library.

- Rosenbaum, D. A. (2007). *MATLAB for Behavioral Scientists*. Mahwah, NJ: Erlbaum. (ISBN 0-8058-6319-2).
- Rosenbaum, D. A. (2010). *Human Motor Control* (2nd ed.). San Diego, CA: Academic Press/Elsevier. [Translated into Japanese, 2012, by MIWA-SHOTEN, LTD, Japan.]
- Rosenbaum, D. A. (2014). *It's a Jungle in There: How Competition and Cooperation in the Brain Shape the Mind*. New York, NY: Oxford University Press.
- Simonsohn, U. (2013). Just post it: The lesson from two cases of fabricated data detected by statistics alone. *Psychological Science*, 24, 1875–1888. doi: 10.1177/0956797613480366.
- Skinner, B. F. (1972). *Cumulative Record*. New York, NY: Appleton-Century Crofts.
- Ulrich, R., & Giray, M. (1989). Time resolution of clocks: Effects on reaction time measurement: Good news for bad clocks. *British Journal of Mathematical and Statistical Psychology*, 42, 1–12.
- Wallisch, P., Lusignan, M. E., Benayoun, M. D., Baker, T. I., Dickey, A. S., & Hasopoulos, N. G. (2009). *MATLAB For Neuroscientists: An Introduction to Scientific Computing in MATLAB*. Burlington, MA: Academic Press/Elsevier.

Commands Index

-	(4.1)	>=	(5.1)
; <i>(output suppression)</i>	(2.4)		(5.1)
; <i>(matrix row delimiter)</i>	(3.1)	~=	(5.1)
: <i>(series delimiter)</i>	(3.2)	abs	(4.2)
...	(2.6)	all	(5.6)
.*	(4.7)	any	(5.6)
./	(4.7)	audioplayer	(11.10)
.^	(4.7)	audioread	(11.13)
' <i>(transpose operator)</i>	(3.5)	audiorecorder	(11.14)
'' <i>(string delimiter)</i>	(6.2)	audiowrite	(11.13)
'' <i>(apostrophe in string)</i>	(6.6)	axis	(9.2)
'bo'	(9.3)	axis equal	(10.4)
'g-'	(9.3)	axis square	(10.2)
's'	(6.6)	bar	(9.16)
()	(4.3)	bar3	(10.7)
[]	(3.8)	barh	(9.16)
{ }	(7.3)	beep	(11.7)
@	(12.4)	Beeper	(13.15)
*	(4.1)	box	(9.10)
* <i>(for matrices)</i>	(4.8)	break	(5.4)
/	(4.1)	brighten	(9.15)
/ <i>(for matrices)</i>	(4.8)	calendar	(2.2)
\n	(6.6)	calendar	(4.11)
\t	(6.6)	camtarget	(10.17)
&	(5.1)	camzoom	(10.17)
%	(2.6)	case	(5.2)
%% <i>(formatting)</i>	(6.6)	cd	(6.10)
%% <i>(section header)</i>	(15.3)	ceil	(4.10)
%d	(6.6)	cell2mat	(7.3)
%e	(6.6)	char	(7.2)
%f	(6.6)	class	(7.1)
%s	(6.6)	clc	(2.2)
^	(4.1)	clear	(3.8)
^ <i>(for matrices)</i>	(4.8)	clear all	(3.8)
+	(4.1)	clf	(9.1)
<	(5.1)	clock	(6.13)
<=	(5.1)	close	(9.1)
==	(5.1)	color	(9.5)
>	(5.1)	colormap	(10.4)

colormap	(9.15)	exist (<i>variable</i>)	(5.1)
comet	(11.2)	exit	(2.2)
comet3	(11.2)	exp	(4.2)
commandwindow	(2.6)	F5 key	(2.6)
compass	(9.14)	fclose	(6.8)
contour	(10.13)	feather	(9.18)
corrcoef	(4.5)	feof	(7.5)
cos	(9.4)	fgetl	(6.15)
cross	(4.8)	fgetl	(7.5)
ctrl-[	(2.6)	figure	(9.1)
ctrl-[	(5.1)	fill	(10.2)
ctrl-]	(2.6)	find	(5.7)
ctrl-]	(5.1)	fix	(4.10)
ctrl-0 (<i>zero</i>)	(2.6)	floor	(4.10)
ctrl-c	(2.2)	fontsize	(10.5)
ctrl-c	(5.4)	fopen	(6.8)
ctrl-i	(2.6)	for	(5.3)
ctrl-i	(5.1)	format	(6.5)
cylinder	(10.16)	format bank	(6.5)
date	(2.2)	format compact	(6.5)
dbclean	(14.1)	format long	(6.5)
dbcont	(14.1)	format long g	(6.5)
dbquit	(14.1)	format loose	(6.5)
dbstep	(14.1)	format rat	(6.5)
deal	(7.4)	format short	(6.5)
delete (<i>handle</i>)	(11.1)	format short g	(6.5)
diary	(2.8)	fprintf	(6.7)
diff	(4.5)	fread	(6.15)
dir	(6.10)	fseek	(6.15)
disp	(2.2)	function	(8.2)
dlmwrite	(6.8)	fwrite	(6.15)
doc	(2.2)	get	(9.5)
dot	(4.8)	get(0, 'ScreenSize')	(9.18)
double	(7.1)	get(gca)	(9.12)
drawnow	(11.3)	get(gcf)	(9.18)
edit	(2.6)	get(h)	(9.10)
else	(5.1)	getframe	(11.4)
elseif	(5.1)	GetMouse	(13.12)
end (<i>case</i>)	(5.2)	GetSecs	(13.2)
end (<i>for</i>)	(5.3)	ginput	(10.5)
end (<i>function</i>)	(8.2)	grid	(9.10)
end (<i>if</i>)	(5.1)	help	(2.2)
end (<i>variable index</i>)	(3.2)	HideCursor	(13.12)
end (<i>while</i>)	(5.4)	hist	(9.15)
errorbar	(9.13)	hold	(9.4)
errordlg	(12.2)	i (<i>imaginary number</i>)	(4.2)
exist	(6.13)	if	(5.1)

image	(10.3)	New Script button	(2.6)
imread	(10.3)	not	(5.1)
input	(6.2)	num2str	(7.2)
inputdlg	(12.2)	ones	(3.6)
iofun	(6.15)	open	(2.2)
isempty	(5.1)	otherwise	(5.2)
isnan	(5.7)	patch	(10.15)
KbCheck	(13.9)	pause	(6.3)
KbName	(13.9)	pie	(9.18)
KbPressWait	(13.14)	play	(11.10)
legend	(9.7)	playblocking	(11.10)
length	(3.4)	plot	(9.1)
light	(10.17)	plot3	(10.8)
linspace	(3.6)	plotyy	(9.18)
listdlg	(12.2)	polar	(9.14)
ListenChar	(13.10)	polyfit	(4.5)
load (.mat file)	(6.14)	polyfit	(9.9)
load (.txt file)	(6.11)	print	(9.17)
log	(4.2)	profile	(15.1)
log10	(4.2)	pwd	(2.2)
log2	(4.2)	pwd	(6.10)
logspace	(3.6)	questdlg	(12.2)
loose	(9.17)	quit	(2.2)
ls	(2.2)	quiver	(9.18)
ls	(6.10)	rand	(4.4)
magic	(4.11)	randi	(4.4)
markeredgecolor	(9.5)	randn	(4.4)
markerfacecolor	(9.5)	randperm	(4.4)
markersize	(9.5)	record.field	(7.4)
max	(4.5)	rem	(4.2)
mean	(4.5)	reshape	(3.5)
median	(4.5)	reshape	(4.4)
mesh	(10.10)	return	(8.2)
meshgrid	(10.9)	rng	(4.4)
min	(4.5)	rotate	(10.17)
mod	(4.2)	rotation	(10.5)
movie	(11.4)	round	(4.10)
movie2avi	(11.5)	Run button	(2.6)
msgbox	(12.2)	save (.mat file)	(6.14)
NaN	(4.6)	saveas	(9.17)
nanmax	(4.6)	sca	(13.3)
nanmean	(4.6)	Screen	(13.3)
nanmedian	(4.6)	Screen ('BlendFunction')	(13.14)
nanmin	(4.6)	Screen ('DrawDots')	(13.13)
nanstd	(4.6)	Screen ('DrawText')	(13.3)
nansum	(4.6)	Screen ('DrawTexture')	(13.7)
nanvar	(4.6)	Screen ('FillOval')	(13.8)

Screen('Flip')	(13.3)	sum	(4.5)
Screen('FrameOval')	(13.6)	surf	(10.11)
Screen('GetFlipInterval')	(13.8)	surfc	(10.14)
Screen('MakeTexture')	(13.7)	surfl	(10.14)
Screen('Openwindow')	(13.3)	switch	(5.2)
Screen('Preference')	(13.5)	text	(9.8)
Screen('TextFont')	(13.5)	TextBounds	(13.15)
Screen('TextSize')	(13.5)	textscan	(6.15)
ScreenTest	(13.2)	textscan	(7.5)
set	(9.12)	tic	(5.5)
set(gca,'YDir','reverse')	(13.12)	timer	(11.10)
set(gca)	(10.1)	title	(9.6)
set(gcf,'Position')	(9.18)	toc	(5.5)
set(gcf)	(10.1)	type	(2.8)
set(h,'Position')	(9.10)	type	(6.8)
SetMouse	(13.12)	uicontrol	(12.3)
SetupPsychtoolbox	(13.2)	uigetdir	(12.2)
shading	(10.17)	uigetfile	(12.2)
shg	(9.1)	uiopen	(12.2)
ShowCursor	(13.12)	up-arrow	(2.5)
sin	(9.1)	var	(4.5)
single	(7.1)	ver	(2.2)
size	(3.4)	VideoReader	(11.6)
sort	(4.9)	VideoWriter	(11.6)
sortrows	(4.9)	view	(10.12)
sound	(11.8)	WaitSecs	(13.3)
soundsc	(11.9)	wavread	(11.11)
sphere	(10.16)	wavwrite	(11.11)
sprintf	(6.6)	while	(5.4)
sqrt	(4.2)	who	(2.2)
stairs	(10.6)	whos	(3.7)
stairs	(9.18)	xlabel	(9.6)
std	(4.5)	xlim	(9.2)
stem	(9.18)	xlsread	(6.12)
str2num	(7.2)	xlswrite	(6.12)
strcat	(7.2)	ylabel	(9.6)
strcmp	(7.5)	ylim	(9.2)
strcmpi	(7.5)	zeros	(3.6)
strfind	(7.5)	zlabel	(10.8)
strrep	(7.5)	zlim	(10.14)
subplot	(9.10)		

Name Index

Allen, W. 4

Baker, T. I. 376

Bempechat, J. 10–11

Benayoun, M. D. 376

Boynton, G. x, 375

Brainard, D. H. 325

Burt, C. 373

Busmeyer, J. R. 373

Cohen, R. G. 256

Dickey, A. S. 376

Diedrich, A. 373

Dweck, C. 10–11

Elliott, M. T. 324

Enserink, M. 374

Farrell, S. 373

Fine, I. x, 375

Fisher, R. 374

Flannery, B. P. 205, 375

Florio, J. 10

Giray, M. 134, 376

Hasopoulos, N. G. 376

Hauser, M. 374

Hayes, B. 375

Heijink, H. 167

James, W. 5

Johnson, R. K. 211–12, 370

Julesz, B. 258

Kamin, L. 373

Kleiner, M. 325

Lewandowsky, S. 373

Lu, C-H. 318, 375

Lusignan, M. E. 376

MacLeod, C. 223

Madan, C. x, 375

Mendel, G. 374

Pelli, D. 325

Plant, R. R. 134, 324

Press, W. H. 205, 375

Proctor, R. 318, 375

Quinlan, P. T. 324

Rock, I. 258

Rosenbaum, D. A. 256

Simonsohn, U. 374

Skinner, B. F. 262, 376

Stapel, D. 374

Stevens, C. 167

Teukolsky, S. A. 375

Tolstoy, L. 4

Turner, G. 134, 324

Ulrich, R. 134, 376

van der Wel, R. 278

Vetterling, W. T. 205, 375

Wallisch, P. x, 376

Walsh, M. 278

Welchman, A. E. 324

Wing, A. M. 324

Subject Index

All MATLAB code appears in `Courier` font, as do all words taken from the code shown in the text body of this book. For a full list of commands see the Commands Index.

- absolute value 65–71
- animation: changing successive images 288–90;
 - comet* command 290; commands 287–8;
 - drawing now 291–2; making movies 291–2;
 - overview 3; practicing 302; reading and running previously saved movies 293–4;
 - saving movies 292
- arrays: assigning literal characters (strings) to variables 136–9; creating, accessing, and using cell arrays 164–8; sorting 90–3; string 136–9
- axes: controlling 212–4; getting and setting properties 232–6; turning on and off 227–30
- bar graphs 239–41; three-dimensional 263–4
- Boolean expression 117, 336
- Boolean operators 100–6, 116
- Boolean values 121, 148, 156–61, 362
- boxes 227–30
- breakpoints 355–62
- built-in utilities 305–7
- calculations 62–98; absolute value 65–71;
 - adding 64–5; calendars 95–6; ceiling 93–5; commands 62–3; correlation 77–9;
 - dividing 64–5; exponentiation 65–71; floor 93–5; generating random numbers 73–7;
 - least-squares fit 77–9; logarithms 65–71;
 - magic squares 95–6; matrix algebra 86–90;
 - maximum 77–9; mean 77–9; minimum 77–9;
 - missing data 79–81; multiplying 64–5;
 - ordering 71–3; overview 3; practicing 96–8;
 - raising values 64–5; remainders 65–71;
 - rounding values 93–5; sorting arrays 90–3;
 - square root 65–71; standard deviation 77–9;
 - subtracting 64–5; sum 77–9; using matrices 81–6; variance 77–9
- calendars, generating 95–6
- ceiling 93–5
- characters: converting 161–4
- central processing unit 5
- Command History window 23
- Command window 164, 185, 195, 214, 241, 256, 295, 305–6, 313, 316, 326, 327, 337, 350–1, 360, 361, 369
- comments (programming) 15, 19, 31, 190;
 - double commenting 371–2
- comparing files 372
- computer architecture 4–6
- contingencies 99–127; applying 121–3;
 - commands 99–100; if-ing instantly 116–19;
 - if-ing instantly and finding indices of satisfying values 119–21; overview 3; practicing 124–7;
 - using `for...end` construct 107–11; using `if...else...end` construct 100–6; using `switch...case...end` construct 106;
 - using `while...end` construct and escaping from run-away loops 111–14; vectorizing rather than using `for...end` 114–16; *see also* Boolean operators
- correlation 77–9
- current directory 24–5, 146–7, 174, 205
- Current Folder window 23, 24, 147
- curves: fitting 225–7
- cylinders 275–7
- data: copying 129–30; meshy 266–7; missing 79–81; overwriting 150–1; pasting 129–30; plotting 236–7; reading data from and writing data to Excel spreadsheets 149; reading data saved as plain text from named files 148–9
- Data Acquisition Toolbox 161
- data types 156–81; applying 178–80; Boolean values 156–61; commands 156; converting characters to numbers and vice versa 161–4; creating, accessing, and using cell arrays 164–8; creating and accessing structures 168–75; identifying strings, numbers, and logical values (Booleans) 156–61; numbers

- 156–61; overview 3; practicing 180–1;
searching and modifying strings 175–8;
strings 156–61
- debugging 355–69; commands 355; error
messages and breakpoints 355–62; graphic
output 364–6; interpreting error messages
363; overview 4; practice 366–69; temporary
feedback 362–3
- desktop tab 32
- diary keeping 35–6
- directories, changing 147
- dividing 64–5
- double commenting 371–2
- Ebbinghaus illusion 284
- editor: finding and navigating 371
- Editor window 23, 31–2, 34, 129, 187, 360,
366, 369, 371
- ellipsoids 277–83
- enhanced user interaction 304–22; built-in
utilities 305–7; commands 304; graphic user
interfaces 305; overview 3; practicing 321–2;
prototyping using GUIDE 312–14; recording
314–21; writing code for user interface
functions 307–12
- error bars: plotting data points with 236–7
- error messages 355–62; interpreting 363
- ethics 373–4
- Excel: reading data from and writing data to
spreadsheets 149
- exponentiation 65–71
- figures: properties 243–4; saving, exporting,
and printing 241–3
- figure windows 260–1
- files, comparing 372
- files, overwriting 150–1
- floor 93–5
- flowcharts 9
- functions 182–207; calling 201–5; commands
182; creating 195–201; drawing on
previously defined versus creating your own
205; input arguments 193–5; multiple inputs
193–5; multiple outputs 192–3; overview
3; practicing 205–7; recursive 201–5; sine
210–2; and subfunctions 197, 200; writing
and using general-purpose functions 187–91
- graphics: output for programs and
debugging 364–6; saving 241–3; *see also*
three-dimensional graphics
- graphic user interfaces 305
- graphs: exporting 241–3; generating other kinds
243–4; printing 241–3; *see also* bar graphs;
plots
- grids 227–30
- GUIDE 312–14; *see also* enhanced user
interaction
- help 24, 28, 54, 306, 329, 370; * 86; axis
253; corrcoef 79; datatypes 159;
dlmwrite 145; factorial 206–7;
format 134; GetSecs 335; ginput 260;
iofun 152; legend 221; logspace 54;
mean 28; or 103; patch 274; pi 28;
plot 216; Screen 329; size 46; view
268; xlsread 149
- Help window 23, 24, 32
- histograms 238–9
- if-ing instantly 116–19; and finding indices of
satisfying values 119–21
- images: adding 332–3; generating 257–60;
loading 255–7; overview 3; practicing 283–6
- input-output 128–55; assigning arrays of literal
characters (strings) to variables 136–9;
checking and changing current directory
146–7; commands 128–9; controlling file
print formats 139–43; copying and pasting
data by hand 129–30; formatting numbers
134–6; getting user input and displaying result
130–3; overview 3; pausing 133; practicing
153–5; precautions against overwriting files
150–1; reaction times and other delays with
tic...toc 133–4; reading data from and
writing data to Excel spreadsheets 149; reading
data saved as plain text from named files
148–9; saving and loading variables in native
MATLAB format 151–2; string arrays 136–9;
writing data to names files 144–5; writing text
to named files 145
- keyboards 5; collecting input 335–5;
monitoring while doing other things 336–7;
responses 304; speed 134; status 347
- Latin square 122
- least-squares fit 77–9
- legends 221–2
- lines: commands 248–9; controlling 214–15;
generating 249–53; having more than one
graph per plot and more types of points and
lines 215–17; overview 3; practicing 283–6
- logarithms 65–71

- magic squares, generating 95–6
- MathWorks 1, 6, 20, 24, 152, 205, 260, 324, 370, 372, 373; GUIDE 312, 313; Parallel Computing 21
- MATLAB 22–37; allowing or suppressing outputs by omitting or including end-of-line semi-colons 29–30; commands 22–3; Command window 24–5; correcting errors in Command window 30–1; keeping a diary 35–6; obtaining and installing 19–20; overview 1–4; practicing interacting 36–7; running and debugging 33–5; website 19; windows 23–4; writing, saving, and running larger programs as scripts 31–3; writing tiny programs in Command window 25–9
- matrices 38–61; clearing and emptying 57–9; commands 38; concatenating 44–6; creating 38–41; emptying 57–9; merge subplots 230–2; overview 3; practicing 59–61; shorthand methods 52–6; size 46–9; specifying elements 41–4; status 56–7; transposing or reshaping 49–52; using 81–6
- matrix algebra 40, 86–90
- MatTap 324
- maximum 77–9
- mean 77–9
- meshgrids 265–6, 270–4
- meshy data 266–7
- minimum 77–9
- missing data 79–81
- modules 182–207; commands 182; multiple outputs from functions 192–3; overview 3; passing multiple input arguments 193–5; practicing 205–7; top-down approach to programming by using modules 182–7; writing and using general-purpose functions 187–91
- mouse data: collecting 339–43
- movies: commands 287; making 291–2; practice 302; reading and running previously saved 293–4; saving 292; *see also* animation; OpenMovie command
- multiplying 64–5
- negative feedback 10–13
- Nelder-Mead simplex direct search algorithm 205
- OCTAVE 20, 325
- OpenMovie command 349–50
- ordering 71–3
- outputs: graphic 364–6; multiple 192–3; programs and debugging 364–6; *see also* input-output
- overwriting files: precautions against 150–1
- parallel processing 5, 6
- pausing 133
- Pearson product-moment correlation coefficient 79
- plots 208–47; adding labels 220–1; adding legends 221–2; adding text 222–5; adding titles 220–1; appearance 214–15; bar graphs 239–41; boxes 227–30; compass 237–8; controlling the appearance of plotted points and lines 214–15; commands 208–10; data points with error bars 236–7; deciding to plot data and generating a sine function 210–2; error bars 236–7; exporting 241–3; fitting curves 225–7; grids 227–30; having more than one graph per plot and more types of points and lines 215–17; histograms 238–9; meshgrids 265–6; meshy data 266–7; multiple graphs 215–17; overview 3; polar 237–8; practicing 244–5; printing 241–3; properties 232–6; subplots 227–30; subplots merger 227–30
- points: controlling 214–15; having more than one graph per plot and more types of points and lines 215–17; plotting 217–20; plotting data points with error bars 236–7; setting 217–20
- polar plots 237–8
- primary memory 5
- print formats: controlling 139–43
- profiling efficiency 372
- programming: being clear about what program should do 8–10; ethically 373–4; need for 7–8; negative feedback 10–11; principles 6–7; productively 370; with a friend 13–14; working incrementally 10
- programs: writing clear 14–17; writing concise 14; writing correct 17–18
- Psychtoolbox 134, 323–54; adding shapes to a display 331–2; adding textures and images to a display 332–3; changing fonts and font sizes 330–1; collecting a response string 337–8; collecting keyboard input 335–6; collecting mouse data 339–43; commands 323–4; creating an animation with moving dots 343–6; displaying stimuli sequentially 333–4; documentation 329–30; installing 325–7; introduction 324–5; making things

- transparent 346–7; monitoring keyboard input 336–7; overview 4, 21; problems 351–4; recovering from crashes and infinite loops 350–1; testing the Simon effect 347–9; writing a simple program 327–9
- raising values 64–5
- random numbers, generating 73–7
- random sequence 76, 77, 113; constrained 121
- reaction times: recording 133–4
- rectangular solids 274–5
- remainders 65–71
- Rotate-3D tool 272
- rounding values 93–5
- saving program scripts and functions 23
- semi-colons 29–30
- serial processing 5, 6
- shapes: filling 253–5; forming 253–5; overview 3; practicing 283–6
- Simon effect 347–9
- sine function 210–2, 301
- sine wave 212, 298–9
- solids 274–5; *see also* three-dimensional graphics
- sorting arrays 90–3
- sound: commands 287–8; controlling volume 295–6; controlling volume while staggering or overlapping 297; delaying 296–7; overlapping 296–7; overview 3; playing beeps 294; practicing 302; related functions 301; staggering or overlapping 296–7
- sound files: creating 297–301; loading and playing 294–5; writing and reading 301
- spheres 275–7
- square root 65–71
- stairing 261–3
- stand-alone applications 372–3
- standard deviation 77–9
- Statistics toolbox 81
- subtracting 64–5
- sum 77–9
- surfing 267–8
- text: adding 222–5
- three-dimensional graphics: bar graphs 263–4; colors 258–60; commands 248–9; contours 269–70; cylinders 275–7; ellipsoids 277–83; meshgrids 265–6, 270–4; plotting 264–5; points of view 268–9; practicing 283–6; rectangular solids 274–5; spheres 275–7
- tic command 115, 127, 134, 314
- titles 220–1
- toc command 115, 127, 134, 314
- user interfaces: prototyping using GUIDE 312–14; recording 314–21; writing code 307–12
- variance 77–9
- vectorizing 114–16
- vectors 39, 54, 86–7, 89–90
- volume, controlling 295–6
- websites: associated with this book 2, 19, 214, 251, 267, 269, 331, 362, 366; MathWorks 21, 152, 373; Psychtoolbox 325, 326, 327
- windows 23–4
- Workspace window 24, 34, 56
- xlabels 220–1
- ylabels 220–1



eBooks

from Taylor & Francis

Helping you to choose the right eBooks for your Library

Add to your library's digital collection today with Taylor & Francis eBooks. We have over 45,000 eBooks in the Humanities, Social Sciences, Behavioural Sciences, Built Environment and Law, from leading imprints, including Routledge, Focal Press and Psychology Press.

ORDER YOUR
FREE
INSTITUTIONAL
TRIAL TODAY

Free Trials Available

We offer free trials to qualifying academic, corporate and government customers.

Choose from a range of subject packages or create your own!

Benefits for you

- Free MARC records
- COUNTER-compliant usage statistics
- Flexible purchase and pricing options
- 70% approx of our eBooks are now DRM-free.

Benefits for your user

- Off-site, anytime access via Athens or referring URL
- Print or copy pages or chapters
- Full content search
- Bookmark, highlight and annotate text
- Access to thousands of pages of quality research at the click of a button.

eCollections

Choose from 20 different subject eCollections, including:

Asian Studies



Economics



Health Studies



Law



Middle East Studies



eFocus

We have 16 cutting-edge interdisciplinary collections, including:

Development Studies



The Environment



Islam



Korea



Urban Studies



For more information, pricing enquiries or to order a free trial, please contact your local sales team:

UK/Rest of World: online.sales@tandf.co.uk

USA/Canada/Latin America: e-reference@taylorandfrancis.com

East/Southeast Asia: martin.jack@tandf.com.sg

India: journalsales@tandfindia.com

www.tandfebooks.com